

Norms of Assertion and Expressivism

Brian Weatherson, Cornell University

In this paper I introduce a new argument for a certain kind of expressivism. The conclusion of the argument is that when we make utterances about the rightness, either moral or prudential, of certain actions, we are not making *assertions*. The premises of the argument are that assertion is governed by a particular norm, the knowledge norm, and this norm does utterances to do with rightness. The argument does not say anything about what we might be doing if not making assertions when we say that something is the right thing to do, but a natural conclusion is that we are expressing our approval of such action.

Although the conclusion of the argument is that, loosely speaking, some form of expressivism is true, the argument doesn't imply that cognitivism is false. At least, it doesn't imply that if cognitivism is the view that claims about rightness have objective, potentially unknowable, truth conditions. Indeed, part of the argument relies on the existence of just such truth conditions. So the argument may not be immediately appealing to those theorists who want to defend expressivism as part of a broader anti-realist project. On the other hand, even they can view the argument as a *reductio* of a kind of strong realist project. 'Strong realism' as I'm using the phrase, is the conjunction of the view that claims about rightness have objective, potentially unknowable, truth conditions, and that one who makes a claim about rightness is asserting the satisfaction of those truth conditions. The argument will put some pressure on that strong realist project.

Two disclaimers before we go any further.

First, the core example that drives this paper is drawn from an (in progress) paper that I'm jointly writing with Ishani Maitra. And the example is basically hers. (And I owe to her many of the clarifications in sections 3 and 4.) I'll take the credit (or blame) for trying to see what expressivist consequences can be drawn from the example however, so this is a sole authored paper. (Moreover, much of the discussion of norms of assertion in section 2 is based on work in progress of hers.)

Second, there is a certain amount of bad faith in my presentation of the argument, since I don't believe its conclusion. I don't think it *should* be a norm that philosophers only present arguments they think are sound (or at least don't believe are unsound) but I suspect that is a norm in some parts. So this is explicit notice that I'm violating any such norm. I think the argument is an interestingly new way to argue for a potentially important conclusion, and that's reason enough to present it. And I'm going to do my best to show why someone could reasonably accept the argument, though I personally don't. (The last section will be about my reasons for doubt.)

While the core conclusion here concerns meta-ethics, the argument will lean heavily on the knowledge norm for assertion, and especially on Timothy Williamson's theory of knowledge. Williamson endorses the knowledge norm, but more importantly his epistemology, in which the notion of safety plays a crucial role, will be central to what follows. So we'll start with a quick tour of the important parts of Williamson's epistemology.

1. Safety and Its Consequences

We'll start with an extreme version of the familiar 'speckled hen' problem for foundationalism. The speckled hen problem is that even when we are looking directly at the hen, we may not know exactly how many speckles it has without counting. (And even if we count, we may lose count.)

We can tell, perhaps, that it has more than 40 and fewer than 60, but that it has 47 exactly is beyond us.

Williamson's version of this example asks us to consider someone trying to judge the size of a crowd at a football stadium. (In US stadiums they sometimes make an interactive game out of this.) Even if everyone is in their seats, and even if we have a full view of the stadium from our vantage point, we will not know exactly how many people are in the stadium. We can make a reasonable estimate of course, but that estimate will have a margin of error to it. So we might know that there are, for example, between 20,000 and 24,000 in the ground, but we won't know that there are exactly 21,974 there, because our crowd estimating skills aren't that good.

What would happen if we guessed, and for some reason came to firmly believe, that there are exactly 21,974 there? Our belief would be true, and it would even be well founded in the evidence. (After all, we believe there are 21,974 there because there are that many there and we are seeing all of them.) But it would not, says Williamson, be knowledge. The reason is that our belief would not be *safe*. A belief is *safe* iff it is not only true, but would be true in all worlds that are (in the relevant sense) similar. There is a lot to be said about what relevant similarity amounts to in this context, but I won't go into that here. Suffice to say that for someone whose crowd judgments have a margin of error, in the ordinary sense, of n , worlds in which the crowd is $21,974 \pm n$ are relevantly similar.

None of this should be particularly controversial. After all, it is a commonplace of statistics that our judgments have a margin of error to them. What is importantly new in Williamson's work is the observation that this undermines a thesis that some epistemologists (and perhaps more epistemic logicians) have found attractive, the so-called KK thesis. The strong form of this thesis is that when S knows that p , then S knows that she knows that p . A weaker form is that when S knows that p , she knows she is in a position to know that p . Williamson uses these facts about margins of error to show that even the weaker thesis is false. I'll illustrate this with two examples.

Billy and Suzy are both looking at our football stadium. Both of them see the 21,974 people, and judge that there are between 20,000 and 24,000 there. In both cases they know this to be true. But in neither case do they know, or are even in a position to know, that it is true.

Billy is not a great judge of crowds. The margin of error on his estimates is about 1500. So when he looks at the crowd, for all he knows there are as few as 20,474 there. (That is, the actual crowd less his margin of error.) Now if there were 20,474 there, he certainly would *not* know that there are more than 20,000 there, because he is not that fine a judge. Indeed, were there 20,474 there, he would not even be in a position to know that there are more than 20,000 present. So although Billy knows that there are more than 20,000 present, he doesn't know that he knows this, or even know that he's in a position to know this.

Suzy is a much better judge than Billy. She can, it turns out, judge crowds to the nearest 500. But she does not know this. For all she knows, she can't judge crowds any more finely than Billy can, i.e. to about the nearest 1500. When Suzy looks at the crowd, for all she knows there are as few as 21,474 there. Moreover, there is an epistemic possibility for Suzy where (a) there are 21,474 there, and (b) her margin of error when measuring crowd sizes is about 1500. So there is an epistemic possibility for her in which she can't tell that the crowd is larger than 20,000. So she does not know that she's even in a position to know the crowd is above 20,000. But she does, indeed, know this.

The point that our knowledge comes apart from what we know we're in a position to know obviously is not restricted to judgments about the size of crowds. It applies whenever there

is a margin of error in our judgments. And that applies everywhere, or at least everywhere that there is an objective fact about which we can make judgments. Hence it applies to moral judgments. We'll look at just such an application, but first we have to look at a connection between knowledge and assertion.

2. Norms of Assertions

Some assertions are better than others. Some, for instance, are funnier than others. Some folk value assertions that are more informative, more salient, shorter, more scandalous, more prurient or what have you. None of these norms are particularly connected to assertion. It's just that various folk value humour, information, salience, brevity, scandal, prurience or what have you. But perhaps there are norms that are distinctively connected to assertion. There are various movements of plastic pieces around a board that have little or no value in themselves, but which are good chess moves. That one is playing chess makes certain actions that would otherwise have no interesting normative properties good or bad. Perhaps the fact that one is making assertions generates new norms that apply to assertions as such, just as certain norms apply to chess moves as such.

Now it isn't obvious that there will be any such norms. Perhaps the rules governing assertion are just the general rules for communal living. But many philosophers have conjectured that this is not true, that there is some norm or other that is distinctively connected to assertion, just like there are some norms that are distinctively connected to various games. In particular, we'll be investigating the idea that the norm *say only what you know*, is distinctively attached to assertion.

There are a few distinctions around here that need cleaning up before we go any further. Some have suggested that the norms attached to assertion, such as the knowledge norm, are to be thought of as like constitutive norms of games. Constitutive rules are what make it the case that one is playing the game that one is actually playing. If a soccer player is off-side, they're breaking the rules, but doing so in a way that is consistent with still playing soccer. If a player picks up the ball and runs forward with it, a la William Webb Ellis, they're playing some other game. Perhaps if one violates a norm of assertion, such as the knowledge norm, then one is doing something else (playing rugby perhaps) rather than asserting.

This suggestion is pretty implausible on the face of it. If it were true, then we'd say that there about as many assertions in the average White House press conference than in a typical performance of *Hamlet*. (Perhaps fewer – one can pretty safely deny that *nothing* is rotten in the state of Denmark.)

But there are other norms in games. A quarterback who throws into quintuple coverage is still playing American football, just playing rather badly. (Or at least recklessly.) The knowledge norm might be construed as being something like that. This seems better; one who asserts what they do not know does seem to be doing badly *qua* asserter. This is most emphatically not to say that they are being simply bad. Sometimes it might be the best thing to do, all things considered, to play a particular game badly. Think of what one should do when playing checkers with a four year old, for example. But the fact that it is good to play checkers badly when playing with four year olds doesn't detract from the fact that such play is *bad checkers play*, and hence that there are norms of checkers that are active and being violated.

Above I was writing as if there is sharp distinction between norms that define the game you're playing, and norms of good play in that game. That probably isn't quite true. Break enough of the norms of good play and you aren't playing the game. Roughly speaking, one of the

norms is *play to win*. So it is constitutive of playing chess that you regard your moves as evaluable by the norms of good chess play. And that's true even if you have reasons to want to make bad chess moves. The same might go for assertion. It might be a constitutive norm of assertion not that you say only what you know, but that you regard your assertions as evaluable by that norm. That, I gather, is something like Williamson's position.

Williamson also seems to be committed to an independent thesis, namely that the speech act of assertion is individuated by the special norm that applies to it. This *isn't* any part of the position I'm defending here. I will be assuming (or perhaps concluding) that the knowledge norm governs assertion, but for all I say here there could be other speech acts that are also governed by just that norm. Whether this claim about individuation is true plays no role in the argument for expressivism that follows, so I'm setting it to one side here.

Why believe that knowledge is a norm of assertion in this sense? There are two primary arguments, an argument from Moore's Paradox and an argument from presupposition. (Williamson also leans heavily on an argument from claims about lotteries. I don't think that works, but the reasons I think it fails are rather involved, and would take us too far from the main thread of this paper.) The Moore's Paradox argument runs as follows.

1. Assertions of the form *p but I don't know that p* are pragmatically defective.
2. The best explanation of the defectiveness is that knowledge is the norm of assertion.
3. So, probably knowledge is the norm of assertion.

We'll see below that premise 1 might be questionable, but not obviously in a way that hurts the expressivist argument. The interest is in premise 2. Assume that whenever I know *p and q* I know *p* and I know *q*. Then I can never know *p but I don't know that p*. To prove this, assume for reductio that I do know it. Then by assumption I know the first conjunct, i.e. *p*. Hence the second conjunct is false. But by assumption I also know the second conjunct. Contradiction. So the knowledge norm does explain the defectiveness of such a claim. This obviously doesn't imply it is the *best* explanation of the phenomena, but it seems pretty clear that no other explanation will be as simple.

The argument from presupposition turns on the following phenomena about conversation. Conversations like the following are *not* defective.

Billy: Alex is happy.

Suzy: How do you know that?

In general, someone can only ask *how do you ϕ ?* if it is *presupposed* that the other person does in fact ϕ . This is quite a strong constraint. Even if A knows that he ϕ 's, and B knows that A ϕ 's, unless A knows that B knows that A ϕ 's, B asking *how do you ϕ ?* may be met with a small amount of shock. (And, notably, by A's asking *how do you know I ϕ ?*) But no such shock will meet Suzy's question. This suggests that when Billy says that *p*, it becomes common ground that *p*, i.e. it becomes presupposed. We then get another abductive inference.

4. When S asserts that *p*, it becomes a presupposition of future conversation that S knows that *p*.
5. The best explanation of this phenomena is that knowledge is the norm of assertion.
6. So, probably, knowledge is the norm of assertion.

The idea behind premise 2 is that it is generally safe to presuppose that people are following rules. E.g., in a game of chess it is not defective to ask after a given move *Why is that a good move?*, which presupposes that it is a good move. Such a question may be *rude*, but that is about the worst that can be said of it. So the knowledge norm does explain the phenomena. Again, we'll set the question of whether it is the *best* explanation to one side. Instead, we'll look at what happens when we try to apply the knowledge norm in cases involving moral talk.

3. An Example concerning Norms of Moral Talk

Our example concerns the country of Indalia. Indalia is a powerful country, and its leaders feel a moral duty to militarily intervene in parts of the world to prevent humanitarian catastrophes. They are currently considering intervening in country X. It is, as it turns out, the right thing to do for them to intervene. (I'm assuming here that such intervention could, *in principle*, be the right thing to do. I'm not assuming that in this respect the case is particularly realistic.)

But it is a close call. Had the evidence they had concerning the impending catastrophe, or of the prospects of success, been a little weaker, it would not have been the right thing to do. And indeed, even with their evidence, had the enemy they are fighting been a little less murderous or the risk of excessive civilian casualties a little higher, it would also have been wrong to go to war. So while it is right to go to war given the evidence, the leaders can't *know* it is right to go to war. The reason is simply that such a belief could not be *safely* true.

So our leaders are in a delicate position here. The Prime Minister of Indalia decides to launch the war, and gives a speech in the House of Commons setting out her reasons. (Note that this is a parliamentary system, where the executive is empowered to launch bombers one day and turn up to Parliament the next day to explain why. That's just what our Prime Minister has done.) She talks about the impending catastrophe, the impossibility of an alternative resolution, and the good prospects of the mission succeeding with light civilian casualties. All of the things she says in the speech are true, and up to her conclusion they are all things that she knows. She concludes with (7).

(7) So, the right thing to do in the circumstances is to go to war.

Now (7) is, we're assuming, true. And, let's add, the Prime Minister believes it. But it is not something she knows. So, if in saying (7) she *asserts* that it is right to go to war, she violates the knowledge norm that governs assertion.

This seems like a crazy result to me. I could just lean on the raw intuition that (7) is a perfectly appropriate thing to say in the circumstances, but I think there are two ways to strengthen the intuition that saying (7) does not violate any norms.

The first point is that if this is a norm violation, then in the story it is the *first* norm violation the Prime Minister has made. It was *right* to send in the bombers, so she violated no norms when she did this. As noted, it was even right to send them in without Parliamentary approval, since she is empowered to launch wars when it is the right thing to do, as it is. Every other assertion she made was something she knew, so there was no violation there. So her first violation was to assert (7). But this is crazy; it can't be wrong to defend in this way a war that you were right to start.

The second point comes from consideration of what she might have said in place of (7). The best she could do would be to recite the evidence supporting the war, and then end with something like (8).

(8) These are all strong reasons for going to war, but I can't say whether they make going to war the right thing to do. Nevertheless, I've sent in the bombers.

I think that (7), which violates the knowledge norm of assertion, is a much more appropriate end to the speech than (8), which does not violate it. And (8) is a better way to end the speech than anything else I can think of that does not violate the norm.

4. *Borderline Wars?*

The above example turns on the possibility of there being borderline cases of right war-making, i.e. actions of war-making that are right but which we cannot know are right. Someone might try and defend the knowledge norm by denying the possibility of such cases. This move obviously wouldn't appeal to Williamson. Indeed, much of what I'll say in response turns on what he has taught concerning safety and the KK principle.

It is very tempting to reason as follows. Let F be a property of circumstances such that the war is objectively just iff the circumstances are, indeed, F . So things might be F iff the enemy is killing its own civilians, if the war won't lead to excessive civilian casualties and so on. Then if circumstances are borderline cases of F ness, the government cannot know that circumstances are F . But a particular agent's action of launching a war is only right if that agent knows that circumstances are F . So there can't be borderline right acts of war-making.

The problem with this line of reasoning is that the most it can show is that any right act of starting a war will start a war that is, objectively speaking, a non-borderline just war. It's going to be important in what follows to distinguish our evaluation of whether from a God's eye perspective we'd judge the war is *just* or not from whether the Prime Minister's action was *right*. I'll exclusively use *just* as an objective (or at least relatively objective) predicate that applies to wars, and *right* as a predicate that applies to particular acts of war-making.

Now it *is* true that some acts of starting just wars are not right. If the Prime Minister starts the war merely because her astrologer said it would work, this is not a *right* action even if the war itself could be defended as *just*. Plausibly, it is *only* right to start a war if one *knows* that the war in question is just. I'll assume that in what follows. (I don't think this tilts the playing field in my favour; indeed, something like it seems to be behind the imagined objection.) In that case, if the war is borderline just, then any belief that the war is just would be unsafe and hence not knowledge, so it would be wrong to start the war. But from this it does not follow that there are no cases where it is borderline *right* to start the war.

The simplest such case comes up in cases where war starting is right given what is known, but some of the crucial evidence consists of facts that we know, but don't know that we know. For example imagine that a crucial piece of the case for war comes from information from a spy working behind enemy lines. As it turns out, the spy is reliable, so we can form knowledge from her testimony. But she could easily enough have been unreliable. She could, for instance, in principle have been bought off by the enemy's agents. As it happens, the amount of money that would have taken was outside the budget the enemy has available for counterintelligence. But had the spy been a little less loyal, or the enemy a little less frugal with the counterintelligence budget, she could easily have been feeding us misinformation. So while the spy is a safe

knowledge source, we don't *know* that she is safe. We don't, for instance, know the size of the enemy's counterintelligence budget, so for all we know she is at risk of being bought off. So if the spy tells us that p , we come to know that p , and p forms part of our evidence. But we don't know that we know that p , so we don't know that p is part of our evidence. So if war starting would not be right if we didn't know that p , we don't know that war starting is right. But that's consistent with us knowing that the war is just, and hence that starting it is right.

A second way in which the Prime Minister's actions could be borderline right is a little trickier. Some wars are borderline just, and others are safely just. But among those that are safely just, some are only barely safely just. They are, in other words, borderline cases of safely just wars. Starting such a war can be right, because we know the war is just. But we can't know that starting the war is right, because that requires knowing that we know the war is just, and that requires the war being safely safely just, not borderline safely just, and it isn't.

That was all rather abstract, and when spelled out in detail it gets a little controversial. You might worry that what I just said conflicts with the supervenience of the normative on the descriptive. That supervenience implies that if the war is just, it is a necessary truth that, given the descriptive facts, the war is just. And you might worry that a war that is necessarily just, given the facts on the ground, can't be borderline just, or even borderline safely just. (Indeed, the way Williamson defines safety sometimes threatens to lead to just this result.)

This is a difficult point to work through all the ramification of, but I think the important point is that we have reason to think that even necessary truths need not be *safely* true. Remember that safety was first introduced as a technical notion to help formalise intuitions about the difference between knowledge and lucky guesses. Now some of the subject matters about which we can have lucky guesses are subjects where all truths are necessarily true, such as mathematics and logic. Given that, we'll want to say that some necessary truths are not safely true. I think the same reasoning suggests that even if it is necessary given the facts on the ground that the war is just, it doesn't follow that the war is safely just. Someone who judges that such a war is just might just be guessing, or at least not reliable in the way required for knowledge, even if the content of their guess is a necessary truth.

5. Expressivism

The above example might look like an argument against the knowledge norm for assertion. But such an argument would have to assume that when the Prime Minister utters (7), she makes an *assertion*. The argument has no force against an expressivist who denies this. Indeed, it can be turned around and used as an argument for the conclusion that uttering (7) is not an assertion.

9. In the above example, the Prime Minister can properly say "It is right to go to war", even if she doesn't know that it is right to go to war.
10. If in saying "It is right to go to war" the Prime Minister asserts that it is right to go to war, then the Prime Minister asserts something she does not know.
11. In all possible situations it is not proper, indeed it is a violation of the knowledge norm, to assert something that one does not know.
12. The example describes a possible situation of a borderline case of right war-making.
13. So the Prime Minister does not assert that it is right to go to war.

As I said at the start, the argument does not say exactly what the Prime Minister *is* doing when she says “It is right to go to war”, though a natural conclusion given that she’s not making an assertion is that she’s expressing her approval of going to war.

In general, we might think that any claim about the rightness of actions (or at least one’s own actions) is proper whenever the action is proper. (The parenthetical restriction is important, and we’ll come back to it below.) The alternative is that it is, in an important respect, *harder* for the world to be such that you can speak to the rightness of ϕ -ing than for the world to be such that you can rightly ϕ . That is, a person who says that it is right to ϕ is (implicitly) making a stronger claim about the state of the world than a person who simply goes ahead and ϕ s is. We seem to have moved into a bizarre world where words speak louder than actions. It is a little tricky to apply principles that are genuinely platitudinous in philosophy, but I think the platitude that actions speak louder than words is meant to rule this possibility out.

If that’s right, then any claim about the rightness of actions (or at least our own actions) is not subject to the knowledge norm. But all assertions are subject to the knowledge norm. So no such claim is an assertion. And that conclusion holds whether the rightness in question is moral rightness or prudential rightness, i.e. whether we are talking about the obligations of morality or merely planning for ourselves. So the expressivist conclusion, although it is in one respect weak because it merely says what a person is *not* doing in making such a claim, is quite broad and in that respect rather strong.

Against all this, it might be claimed that it is never right to make assertions that are mere guesses. But if all that is required for a claim about rightness to be proper is that it be true, then some guesses will be correct assertions. The response to this objection is to note that for an important class of rightness claims, a certain kind of mere guess *cannot* be true. Picking out this class requires some care. A claim *S is right to ϕ* falls into this class if it meets two conditions.

- S is (known to be) the speaker, or it is known that it is right for S to ϕ iff it is right for the speaker to ϕ if they were in S’s circumstances
- ϕ -ing requires positive reasons; it is never right to ϕ merely because there is no reason not to ϕ .

Note that I’m not distinguishing between moral rightness and prudential rightness explicitly here, though the second disjunct of the first clause will be easier to satisfy in cases where the rightness in question is moral rather than prudential.

If those two conditions are met, then in *one sense* the speaker can’t truly say *S is right to ϕ* if they are just guessing. That’s because by the first clause their claim is only true if it is right for them to ϕ in the circumstances S is in. And the second clause implies that they have a reason to ϕ . Now it might be, as in our primary example, that the speaker doesn’t know that the reason is a sufficient reason even though it is in fact a sufficient reason. And in *that sense* they might be guessing; they might be guessing that a particular reason to ϕ is a sufficient reason in the circumstances to ϕ . But they cannot truly believe this for no reason at all, and so in a particularly strong sense of guess, they can’t truly guess that such a claim is true.

Not coincidentally, I think the classes of rightness claims meeting these two conditions is the class of claims to which the argument applies. (That’s why I had the earlier parenthetical asides about restriction to my own actions.) For all that is said in the argument, it is possible that some rightness claims really are assertions, not expressions. For a case where this seems particularly plausible, imagine that I’m having a casual discussion about mergers and acquisitions

in the tech industry with some friends at the bar. When I say something like *It is right for MegaTech to pursue the merger with BioQuest*, it seems plausible that I'm asserting something about the utility (or profit) maximising action for MegaTech. Since I couldn't possibly be in MegaTech's position, I'm not expressing any attitude about what to do in such a position.

It might be worried that there is a Frege-Geach style worry around here. I've said that there might be a striking difference between the following two utterances.

(14) It is right for me to ϕ .

(15) It is right for S to ϕ .

Compositionality considerations, and more general plausibility considerations, seem to tell against this. The force of such considerations is strongest when we think of cases of mistaken identity, where S is the speaker even though the speaker does not realise this. It is very hard to imagine that the truth conditions of (14) and (15) could differ in these cases. But the argument here does *not* suggest that the truth conditions of (14) and (15) are different. It merely suggests that what a speaker *does* is different in the two cases. So some of the problems besetting traditional expressivism won't be problems for the kind of expressivism suggested by this argument.

6. Moore's Paradox

I'm going to conclude by considering the two arguments for the knowledge norm of assertion, and considering whether they give us reason to overturn our judgment about the core example. The conclusion will be that they don't, but there is one worry for the expressivist argument. The worry is that these arguments may fail so badly that we are forced to give up the knowledge norm for assertion. In that case we won't have an argument left that what the Prime Minister does is not an assertion. I'll start with the good news for expressivists, and then move on to the bad news.

As noted at the top, one of the core arguments for the knowledge norm is that if knowledge is not a norm of assertion, then certain Moore paradoxical sentences would be acceptable. But Moore paradoxical sentences are never acceptable. Hence knowledge is a norm of assertion. So if Moore paradoxical sentences concerning normative matters are defective, there is an argument that the knowledge norm applies to them. The simple response to this argument is that it really isn't obvious that Moore paradoxical sentences in these cases are defective. To see this, imagine the Prime Minister in our little story being interviewed on *Newsnight*.

Interviewer: Is it right to launch this war?

Prime Minister: Yes it is. The enemy is causing a humanitarian crisis, they are killing their own civilians as we speak, and our military has the power to stop it.

Interviewer: But wars always carry risks. And some people in the Commons disagree with you. Do you know this is the right thing to do?

Prime Minister: Well you almost never *know* that launching a war is the right thing to do, and I certainly don't in this case. That's one of the hard things about leadership, often times you are faced with hard choices where whatever you decide, you won't know that what you're doing is right. So I'm not surprised there is disagreement, because this is one of those times, and we have to rely on our judgment, not just on our knowledge. But going to war, defeating this tyrant, *is* the right thing to do, and when we look back at this war with the benefit of hindsight, everyone will agree we did the right thing.

That little speech of the Prime Minister's might be somewhat self-serving, even a little sanctimonious, and it isn't obviously a vote winner. But it is at least a little more gracious to her political opponents than some speeches that real world leaders have made in similar situations, because the Prime Minister acknowledges that she doesn't know, that no one in her position could know, whether she's doing the right thing by going to war. Maybe I'm just being too generous because of this graciousness, but it seems to me that the Prime Minister's little speech is not thereby defective in the way that Moore paradoxical utterances are alleged to be always defective. And that's so even though she explicitly says that going to war is the right thing to do, and that she doesn't know this is true.

So there's a fairly flat footed answer to the Moorean argument, namely that in these cases it *isn't* defective to utter a Moore paradoxical sentence. It might be thought that this undermines the original argument for the knowledge norm and hence the argument for expressivism. The expressivist can make one defensive move here, and one offensive move. The defensive move is that in cases where the utterance is uncontroversially an *assertion*, Moore paradoxicality is still defective, and hence her argument is not undercut. The offensive move is to note that this is a rather striking distinction between descriptive and normative sentences, in the former case we can never say *p but I don't know that p*, in the latter case we can, and this distinction needs explaining. Of course, the expressivist has an explanation ready to hand...

7. Norms and Presupposition

A deeper problem for the expressivist argument comes from thinking more closely about the connection between norms and presuppositions. Above I used the argument that knowledge must be a norm of assertion because otherwise hearers wouldn't be entitled to presuppose that speakers know the truth of their utterances. The general principle is that if it becomes common ground in a conversation that one of the participants ϕ 'd, and proper ϕ -ing requires doing X, other participants are entitled to presuppose that she did in fact do X.

The worry obviously is that we only get from the principle and the fact about presupposition to the knowledge norm by a rather unsubtle instance of affirming the consequent. If there are *other* ways to generate presuppositions, the knowledge norm would not be supported by the data. And I suspect that thinking about more mundane cases suggests such ways.

Consider the following two dialogues.

- (16) Suzy: I got a new laptop computer.
Billy: When did you buy it?
- (17) Suzy: I got a new laptop computer
Billy: ??When did you steal it?

Unless Suzy and Billy regularly engage in criminal activity, Billy's reply in (17) is crashingly bad in several respects. His response in (16) is generally good. You might think that this helps the knowledge norm argument. After all, his response in (16) presupposes that Suzy *bought* her computer. And the explanation for this is that the normatively appropriate way to get a new computer is to buy it, and Billy is entitled to presuppose that things go in the normatively appropriate way, absent evidence to the contrary.

But that explanation can't be quite right. For there are other normatively appropriate ways to get a new computer. Suzy might have been *given* the computer. Indeed, if Billy and Suzy

are in a community where one is regularly given such things (perhaps they are rich undergraduates, or faculty at departments with generous research funds), then Billy's question starts to sound a little defective. The right principle seems to be something more like this.

- If it becomes common ground in a conversation that one of the participants ϕ 'd, and proper ϕ -ing usually involves doing X, other participants are entitled to presuppose that she did in fact do X.

Now apply this back to knowledge and assertion. There are all sorts of alternatives to the knowledge norm that are slightly weaker. For instance, proper assertion may require justified belief, or perhaps it requires having undefeated reasons to believe. (One can motivate these norms a little by thinking about Gettier cases, though intuitions there are pretty hazy.) It is hard to *plan* to have a justified belief that doesn't amount to knowledge, or to have undefeated reasons to believe that p without knowing that p . So in either case, the standard way to satisfy the norm of assertion for p is to know that p . So given our principle, even if the norm of assertion is the justified belief norm, or the undefeated reasons norm, other people would be entitled to presuppose that the speaker knows that p , just as Billy is entitled to presuppose that Suzy bought her new computer.

This all is a problem for the expressivist argument I'm offering. Let's think in particular about the undefeated reasons norm. If our Prime Minister is right to go to war, so her reasons for going to war are undefeated, then it is very plausible that her reasons for judging it right to go to war are undefeated. (There is an interesting kind of transparency claim here. Her reasons for judging it right to go to war just are her reasons to go to war. The rightness is 'transparent'. There is an analogy to be drawn with transparency claims about perceptual sensation, but I won't labour that analogy *here*.) So she is right to *assert* that it is right to go to war, even though she does not know it. And this explains why she can make Moore paradoxical utterances; she also has undefeated reason to believe that she does not know. In general it is very hard to have undefeated reason to believe that p and undefeated reason to believe that you don't know that p . What is distinctive about the class of rightness claims at issue here is, if my imagined assertivist defence works, that this isn't as hard when p is such a rightness claim.

So the expressivist argument might not be totally compelling. But at the very least the anti-expressivist has to do a bit of work to reject it, perhaps giving up principles that she'd otherwise quite like to accept about knowledge and assertion.