

Assertion, Knowledge and Action

Ishani Maitra, Brian Weatherson

It is widely believed that the mere truth of p is insufficient for p to be properly assertable, even if p is relevant to current conversation. If a speaker simply guessed that p is true, then she shouldn't say p , for example. There is some dissent from this view (e.g., Weiner (2005)), but it is something close to orthodoxy in the current literature on assertion that something further is needed. The most common 'something else' is *knowledge*: a speaker shouldn't say p unless they know p . This view is nowadays commonly associated with Timothy Williamson (1996; 2000), but it has historical antecedents tracing back at least to Max Black's 1952 paper "Saying and Disbelieving".¹ Call Williamson's position *The Knowledge Rule*.

The Knowledge Rule Assert that p only if you know that p .

This paper aims to raise trouble for The Knowledge Rule, and several related positions, by focussing on a particular kind of assertion. We'll be looking at assertions about what is to be done. The boldest statement of our position is that if an agent should do X , then that agent is in a position to say that they should do X . (We'll qualify this a little below, but it's helpful to start with the bold position.) We argue, following Williamson's 'anti-luminosity' arguments, that its being true that X is the thing to do for an agent doesn't entail that that agent knows it's the thing to do.² If both these claims are true, then there will be cases where it is fine to assert that X is what to do, even though the agent doesn't know this. So, The Knowledge Rule is mistaken. Slightly more formally, we'll be interested in arguments of this structure.

[†] Penultimate draft only. Please cite published version if possible. Final version forthcoming in *Philosophical Studies*. We'd like to thank Matthew Benton, Jessica Brown, Andy Egan, and Susanna Schellenberg, as well as an audience at the Bellingham Summer Philosophy Conference, for helpful discussion of earlier drafts of this paper.

¹ Timothy Williamson (Williamson, 2000, Ch. 11) has the clearest statement of the view we're considering here. It is also defended by Keith DeRose (2002). Both DeRose and John Hawthorne (2004) deploy it extensively as a constraint on theories of knowledge. Jason Stanley (2008) argues for an even stronger constraint: that we should only assert p if we are certain that p . Igor Douven (2006) argues that truth is neither sufficient nor necessary, so the norm should be assert only what is rationally credible. Kent Bach (2010) and Frank Hindriks (2007) both suggest that the only real norm governing *assertion* is belief, but that since knowledge is a norm of belief, we shouldn't generally assert what we do not know. In this paper we're not concerned with the question of whether the rule *Assert only what you know* holds solely in virtue of the normative nature of assertion itself, as Williamson thinks, or partly in virtue of norms applying to related states like belief, as Bach and Hindriks suggest, but rather whether the rule is even a good rule.

² We'll use the expressions 'thing to do' and 'what to do' interchangeably throughout the paper. By X is *what to do*, we mean X ought to be done, all things considered. We take no position on whether X 's being what to do entails its being the morally right thing to do. That may be the case, but nothing we say in this paper depends on its being so.

Master Argument (First Attempt)

1. If act X is what to do for agent S , then S can properly assert that X is what to do (assuming that this assertion is relevant to the current conversation).
2. It is possible that X is what to do for S , even though S is not in a position to know this.
3. So, it is possible that S can properly assert that X is what to do even though she does not know, and is not even in a position to know, that X is what to do.

In section 1, we'll motivate premise 1 with a couple of vignettes. In section 2, we'll qualify that premise and make it more plausible. In section 3, we'll motivate premise 2. In section 4, we'll look at one of the positive arguments for The Knowledge Rule, the argument from Moore's paradox, and conclude that it is of no help. In section 5, we'll look at what could be put in place of The Knowledge Rule, and suggest two alternatives.

The Evidence Responsiveness Rule Assert that p only if your attitude towards p is properly responsive to the evidence you have that bears on p .

The Action Rule Assert that p only if acting as if p is true is the thing for you to do.

We're not going to argue for these rules in detail; that would take a much longer paper. Nor are we going to decide between them. What we are going to suggest is that these rules have the virtues that are commonly claimed for The Knowledge Rule, but lack The Knowledge Rule's problematic consequences when it comes to assertions about what to do.

1 Speaking about What to Do

We start by motivating premise 1 of the Master Argument with a couple of examples. Both cases are direct counterexamples to The Knowledge Rule, but we're interested in the first instance in what the cases have in common. After presenting the vignettes, we offer three distinct arguments to show that, in such cases, it *is* proper for the speakers to assert what they do assert, even though they don't know it to be true.

Going to War

Imagine that a country, Indalia, finds itself in a situation in which the thing for it to do, given the evidence available to its leaders, is to go to war against an enemy. (Those pacifists who think it is never right to go to war won't like this example, but we think war can at least sometimes be justified.) But it is a close call. Had the evidence been a bit weaker, had the enemy been a little less murderous, or the risk of excessive civilian casualties a little higher, it would have been preferable to wait for more evidence, or use non-military measures to persuade the enemy to change its ways. So, while going to war is the thing to do, the leaders of Indalia can't know this. We'll come back to

this in section 2, but the crucial point here is that knowledge has a safety constraint, and any putative knowledge here would violate this constraint.

Our leaders are thus 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. All the things she says in the speech are true, and up to her conclusion they are all things that she knows. She concludes with (1).

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

Now (1) is also true, and the Prime Minister believes it, but it is not something she knows. So, the Prime Minister violates The Knowledge Rule when she asserts (1). But it seems to us that she doesn't violate any norms in making this assertion. We'll have a lot more to say about why this is so in a few paragraphs. But first, here's a less dramatic case that is also a counterexample to The Knowledge Rule, one that involves prudential judgments rather than moral judgments.

Buying Flood Insurance

Raj and Nik are starting a small business. The business is near a river that hasn't flooded in recent memory, but around which there isn't much flood protection. They could buy flood insurance which would be useful in a flood, naturally, but would be costly in the much more likely event that there is not a flood. Raj has done the calculations of the likelihood of a flood, the amount this would damage the business, the utility loss of not having this damage insured, and the utility loss of paying flood insurance premiums. He has concluded that buying flood insurance is the thing to do. As it happens, this was a good conclusion to draw: it does, in fact, maximise his (and Nik's) expected utility over time. (It doesn't maximise their actual utility, as there actually won't be a flood over the next twelve months. So, the insurance premium is an expense they could have avoided. But that doesn't seem particularly relevant for prudential evaluation. Prudential buyers of insurance should maximise expected utility, not actual utility. Or so we must say unless we want to be committed to the view that everyone who buys an insurance policy and doesn't make a claim on it is imprudent.)

But again, it's a close call. If there had been a little less evidence that a flood was a realistic possibility, or the opportunity cost of using those dollars on insurance premiums had been a little higher, or the utility function over different outcomes a little different, it would have been better to forego flood insurance. That suggests that safety considerations make it the case that Raj doesn't know that buying flood insurance is the thing to do, though in fact it is.

Let's now assume Raj has done everything he should do to investigate the costs and benefits of flood insurance. We can imagine a conversation between him and Nik going as follows.

Nik: Should we get flood insurance?

Raj: I don't know. Hold on; I'm on the phone.

Nik: Who are you calling?

Raj: The insurance agent. I'm buying flood insurance.

There is clearly a pragmatic tension in Raj's actions here. But given The Knowledge Rule, there's little else he can do. It would be a serious norm violation to say nothing in response to Nik's question. And given that he can't say "Yes" without violating The Knowledge Rule, he has to say "I don't know". Moreover, since by hypothesis buying flood insurance is the thing to do in his situation, he can't not buy the insurance without doing the wrong thing. So, given The Knowledge Rule, he's doing the best he can. But it's crazy to think that this is the best he can do.

We think that these cases are problems for The Knowledge Rule. In particular, we think that in each case, there is a non-defective assertion of something that is not known. It seems to us intuitively clear that those assertions are non-defective, but for those who don't share this intuition, we have three independent arguments. The arguments focus on **Going to War**, but they generalize easily enough to **Buying Flood Insurance**.

Argument One: "That was your first mistake"

Imagine that the Prime Minister has a philosophical advisor. And the advisor's job is to inform the Prime Minister whenever she violates a norm, and stay silent otherwise. If The Knowledge Rule is correct, then the advisor should stay silent as the Prime Minister orders the invasion, silent as the Prime Minister sets out the reasons for the invasion, then speak up at the very last line of the speech. That strikes us as absurd. It's particularly absurd when you consider that the last line of the speech is supported by what came earlier in the speech, and the Prime Minister believes it, and asserts it, because it is well supported by what came earlier in the speech. Since we think this couldn't be the right behaviour for the advisor, we conclude that there's no norm violation in the Prime Minister asserting (1).

We've heard two replies to this kind of argument. According to one sort of reply, The Knowledge Rule is not meant to be an 'all-things-considered' norm. The defender of The Knowledge Rule can say that the Prime Minister's assertion is defective because it violates that rule, but allow that it is nevertheless all-things-considered proper, because some other norm outweighs The Knowledge Rule on this occasion. We agree that The Knowledge Rule is not intended to be an all-things-considered norm. But even keeping clearly in mind the distinction between being defective in some respect and being defective all-things-considered, it is still deeply unintuitive to say that the Prime Minister's assertion is defective in a respect. That is, we don't think the philosophical advisor should speak up just at the very end of the Prime Minister's speech even if she's meant to observe *all* the norm violations (rather than just the all-things-considered norm violations).

Perhaps the defender of The Knowledge Rule needn't just appeal to an intuition here. Another reply we've heard starts from the premise that the Prime Minister's assertion would be better, in a certain respect, if she *knew* that it was true. Therefore, there is a respect in which that assertion is defective, just as The Knowledge Rule requires. To this second reply, our response is that the premise is true, but the reasoning is invalid. Saying why requires reflecting a bit on the nature of norms.

There are lots of ways for assertions to be better. It is better, *ceteris paribus*, for assertions to be funny rather than unfunny. It is better for assertions to be sensitive

rather than insensitive. (We mean this both in the Nozickian sense, i.e., an assertion is sensitive iff it wouldn't have been made if it weren't true, and in the Hallmark greeting card sense.) It is better for speakers to be certain of the truth of their assertions than for them to be uncertain. But these facts don't imply that humour, sensitivity, or certainty are norms of assertion, for it doesn't follow that assertions that lack humour (or sensitivity or certainty) are *always* defective. Similarly, the fact that it is better to know what you say than not doesn't imply that asserting what you don't know is always defective. In slogan form: *Not every absence of virtue is a vice*. We think knowledge is a virtue of assertions. (In fact, we think that pretty much every norm of assertion that has been proposed in the literature picks out a virtue of assertion.) What we deny is that the absence of knowledge is (always) a vice. Since not every absence of virtue is a vice, one can't argue that the Prime Minister's assertion is *defective* by arguing it could have been better. And that's why the argument being considered is invalid.

Argument Two: "Actions speak louder than words"

It's a bit of folk wisdom that actions speak louder than words. It isn't crystal clear just what this wisdom amounts to, but we think one aspect of it is that an agent incurs more normative commitments by *doing X* than by *talking* about X. But if The Knowledge Rule is right, then this piece of wisdom is in this aspect back-to-front. According to that rule, an agent incurs a greater normative commitment by *saying* that X is what to do than they do by just doing X. If they do X, and X is indeed what to do, then they've satisfied all of their normative commitments. If, by contrast, they say that X is what to do, then not only must X be what to do, but they must know this fact as well. This strikes us as completely back-to-front. We conclude that there is nothing improper about asserting that X is what to do (as the Prime Minister does), when X is in fact what to do.

Argument Three: "What else could I do?"

Here's a quite different argument that *Going to War* is a counterexample to The Knowledge Rule.

1. If ending the speech the way she did was a norm violation, there is a better way for the Prime Minister to end her speech.
2. There is no better way for the Prime Minister to end the speech without saying something that she does not know to be true.
3. So, ending the speech the way she did was not a norm violation.
4. So, The Knowledge Rule is subject to counterexample.

Premise 1 is a kind of 'ought-implies-can' principle, and as such, it isn't completely obvious that it is true. But when we've presented this argument to various groups, the focus has always been on premise two. The common complaint has been that the Prime Minister could have ended the speech in one of the following ways, thereby complying with The Knowledge Rule.

- I've decided that going to war is the thing to do in the circumstances.

- I believe that going to war is the thing to do in the circumstances.
- It seems to me that going to war is the thing to do in the circumstances.

Our first reply to this suggestion is that we'd fire a speechwriter who recommended that a Prime Minister end such a speech in such a weaselly way, so this hardly counts as a criticism of premise 2. Our more serious reply is that even if the Prime Minister ended the speech this way, she'd *still* violate The Knowledge Rule. To see why this is so, we need to pay a little closer attention to what The Knowledge Rule says.

Note that The Knowledge Rule is *not* a rule about what kind of declarative utterance you can properly make. An actor playing Hamlet does not violate The Knowledge Rule if he fails to check, before entering the stage, whether something is indeed rotten in the state of Denmark. The rule is a rule about what one *asserts*. And just as you can assert *less* than you declaratively utter (e.g., on stage), you can also assert *more* than you declaratively utter.³ For instance, someone who utters *The F is G* in a context in which it is common ground that *a* is the *F* typically asserts both that the *F* is *G*, and that *a* is *G*. Similarly, someone who utters *I think that S* typically asserts both asserts that they have a certain thought, and asserts the content of that thought. We can see this is so by noting that we can properly challenge an utterance of *I think that S* by providing reasons that *S* is false, even if these are not reasons that show that the speaker does not (or at least did not) have such a thought. In the context of her speech of the House of Commons, even if the Prime Minister were to end with one of the options above, she would still assert the same thing she would assert by uttering (1) in the circumstances, and she'd still be right to make such an assertion.

2 Bases for Action and Assertion

One might worry that premise 1 in our master argument is mistaken, in the following way. We said that if *X* is the thing to do for *S*, then *S* can say that *X* is what to do. But one might worry about cases where *S* makes a lucky guess about what is to be done. Above we imagined that Raj had taken all of the factors relevant to buying flood insurance into account. But imagine a different case, one involving Raj*, Raj's twin in a similar possible world. Raj* decides to buy flood insurance because he consults his Magic 8-Ball. Then, even if buying flood insurance would still maximize his expected utility, it doesn't seem right for Raj* to say that buying flood insurance is what to do.

Here is a defence of premise 1 that seems initially attractive, though not, we think, ultimately successful. The Magic 8-ball case isn't a clear counterexample to premise 1, it might be argued, because it isn't clear that buying flood insurance for these reasons is the thing for Raj* to do. On one hand, we do have the concept of doing the right thing for the wrong reasons, and maybe that is the right way to describe what Raj* does if he follows the ball's advice. But it isn't clearly a correct way to describe Raj*. It's not true, after all, that he's maximising actual utility. (Remember that there will

³The points we're about to make are fairly familiar by now, but for more detail, see Cappelen and Lepore (2005), which played an important role in reminding the philosophy of language community of their significance.

be no claims on the policy he buys.) And it isn't clear how to think about expected utility maximisation when the entrepreneur in question relies on the old Magic 8-Ball for decision making. And we certainly want to say that there's something wrong about this very decision when made using the Magic 8-Ball. So, perhaps we could say that buying flood insurance isn't what to do for Raj* in this variant example, because he has bad reasons.

But this seems like a tendentious defence of the first premise. Worse still, it is an unnecessary defence. What we really want to focus on are cases where people do the right thing for the right reasons. Borrowing a leaf from modern epistemology, we'll talk about actions having a basis. As well as there being a thing to do in the circumstances (or, more plausibly, a range of things to do), there is also a correct *basis* for doing that thing (or, more plausibly, a range of correct bases). What we care about is when *S* does *X* on basis *B*, and doing *X* on basis *B* is the thing to do in *S*'s situation. Using this notion of a basis for action, we can restate the main argument.

Master Argument (Corrected)

1. If doing *X* on basis *B* is what to do for agent *S*, then *S* can properly, on basis *B*, assert that *X* is what to do (assuming this is relevant to the conversation).
2. It is possible that doing *X* on basis *B* is what to do for *S*, even though *S* is not in a position to know, and certainly not in a position to know on basis *B*, that *X* is what to do.
3. So, it is possible that *S* properly can assert that *X* is what to do, even though she does not know, and is not even in a position to know, that *X* is what to do.

We endorse this version of the master argument. Since its conclusion is the denial of The Knowledge Rule, we conclude that The Knowledge Rule is mistaken. But we perhaps haven't said enough about premise 2 to seal the argument. The next section addresses that issue.

3 Marginal Wars

The argument for premise 2 is just a simple application of Williamson's anti-luminosity reasoning. (The canonical statement of this reasoning is in (Williamson, 2000, Ch. 4)).) Williamson essentially argues as follows, for many different values of *p*. There are many ways for *p* to be true, and many ways for it to be false. Some of the ways in which *p* can be true are extremely similar to ways in which it can be false. If one of those ways is the actual way in which *p* is true, then to know that *p* we have to know that situations very similar to the actual situation do not obtain. But in general we can't know that. So, some of the ways in which *p* can be true are not compatible with our knowing that *p* is true. In Williamson's nice phrase, *p* isn't *luminous*, where a luminous proposition is one that can be known (by a salient agent) whenever it is true. The argument of this paragraph is called 'an anti-luminosity argument', and we think that many instances of it are sound.

There is a crucial epistemic premise in the middle of that argument: that we can't know something if it is false in similar situations. There are two ways that we could try to motivate this premise. First, we could try to motivate it with the help of conceptual considerations about the nature of knowledge. That's the approach that Williamson takes. But his approach is controversial. It is criticised by Sainsbury (1996) and Weatherson (2004) on the grounds that his safety principle goes awry in some special cases. Sainsbury focuses on mathematical knowledge, Weatherson on introspective knowledge. But the cases in which we're most interested in this paper – Indalia going to war, Raj and Nik buying flood insurance – don't seem to fall into either of these problem categories. Nevertheless, rather than pursue this line, we'll consider a different approach to motivating this premise.

The second motivation for the epistemic premise comes from details of the particular cases. In the two cases on which we're focusing, the agents simply lack fine discriminatory capacities. They can't tell some possibilities apart from nearby possibilities. That is, they can't know whether they're in one world or in some nearby world. That's not because it's conceptually impossible to know something that fine, but simply an unfortunate fact about their setup. If they can't know that they're not in a particular nearby world in which $\neg p$, they can't know p . Using variants of *Going to War*, we'll describe a few ways this could come about.

The simplest way for this to come about is if war-making is the thing to do given what we know, but some of the crucial evidence consists of facts that we know, but don't know that we know. Imagine that a crucial piece of Indalia's case for war comes from information from an Indalian spy working behind enemy lines. As it turns out, the spy is reliable, so the leaders of Indalia can acquire knowledge from her testimony. But she could easily enough have been unreliable. She could, for instance, 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 counter-intelligence. 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 supplying misinformation to Indalia. So, while the spy is a safe knowledge source, the Indalian leaders don't *know* that she is safe. They don't, for instance, know the size of the enemy's counterintelligence budget, or how much it would take to buy off their spy, so for all they know, she is very much at risk of being bought off.

In this case, if the spy tells the Indalian leaders that p , they come to know that p , and they can discriminate p worlds from $\neg p$ worlds. But they don't know that they know that p , so for all they know, they don't know p . And for some p that they learn from the spy, if they don't know p , then going to war isn't the thing for them to do in the circumstances. So, given that they don't know the spy is reliable, they don't know that going to war is the thing for them to do. But the spy really is reliable, so they do know p , so going to war is indeed the thing for them to do.

Or consider a slightly less fanciful case, involving statistical sampling. Part of the Prime Minister's case for starting the war was that the enemy was killing his own citizens. Presumably she meant that he was killing them in large numbers. (Every country with capital punishment kills its own citizens, but arguably that isn't a sufficient reason to invade.) In practice, our knowledge of the scope of this kind of

governmental killing comes from statistical sampling. And this sampling has a margin of error. Now imagine that the Indalian leaders know that a sample has been taken, and that it shows that the enemy has killed n of his citizens, with a margin of error of m . So, assuming there really are n killings, they know that the enemy has killed between $n - m$ and $n + m$ of his citizens. Since knowing that he's killed $n - m$ people is sufficient to make going to war the thing to do, the war can be properly started.

But now let's think about what the Indalian leaders know that they know in this case. The world where the enemy has killed $n - m$ people is consistent with their knowledge. And their margin of error on estimates of how many the enemy has killed is m . So, if that world is actual, they don't know the enemy has killed more than $n - 2m$ of his citizens. And that knowledge might not be enough to make going to war the thing to do, especially if m is large. (Think about the case where $m = n/2$, for instance.) So, there's a world consistent with their knowledge (the $n - m$ killings world), in which they don't know enough about what the enemy is doing to make going to war the thing to do. In general, if there's a world consistent with your knowledge where p is false, you don't know p . Letting p be *Going to war is what to do*, it follows then that they don't know that going to war is what to do, even though it actually is the thing to do.

Another way we could have a borderline war is a little more controversial. Imagine a case where the leaders of Indalia know all the salient descriptive facts about the war. They know, at least well enough for present purposes, what the costs and benefits of the war might be. But it is a close call whether the war is the thing to do given those costs and benefits. Perhaps different plausible moral theories lead to different conclusions. Or perhaps the leaders know what the true moral theory is, but that theory offers ambiguous advice. We can imagine a continuum of cases where the true theory says war is clearly what to do at one end, clearly not what to do at another, and a lot of murky space between. Unless we are willing to give up on classical logic, we must think that somewhere there is a boundary between the cases where it is and isn't what to do, and it seems in cases near the boundary even a true belief about what to do will be unsafe. That is, even a true belief will be based on capacities that can't reliably discriminate situations where going to war is what to do from cases where it isn't.

We've found, when discussing this case with others, that some people find this outcome quite intolerable. They think that there must be some epistemic constraints on war-making. And we agree. They go on to think that these constraints will be incompatible with the kind of cases we have in mind that make premise 2 true. And here we disagree. It's worth going through the details here, because they tell us quite a bit about the nature of epistemic constraints on action.

Consider all principles of the form

(KW) Going to war is $N1$ only if the war-maker knows that going to war is $N2$.

where $N1$ and $N2$ are normative statuses, such as being the thing to do, being right, being good, being just, being utility increasing, and so on. All such principles look

like epistemic constraints on war-making, broadly construed. One principle of this form would be that going to war is right only if the war-maker knows that going to war is just. That would be an epistemic constraint on war-making, and a plausible one. Another principle of this form would be that going to war is the thing to do only if the war-maker knows that going to war increases actual utility. That would be a very strong epistemic constraint on war-making, one that would rule out pretty much every actual war, and one that is consistent with the anti-luminosity argument with which we started this section. So, the anti-luminosity argument is consistent with there being quite strong epistemic constraints on war-making.

What the anti-luminosity argument is *not* consistent with is there being any true principle of the form (KW) where $N1$ equals $N2$. In particular, it isn't consistent with the principle that going to war is the thing to do only if the war maker knows that it is the thing to do. But that principle seems quite implausible, because of cases where going to war is, but only barely, the thing to do. More generally, the following luminosity of action principle seems wrong for just about every value of X .

(LA) X is the thing for S to do only if S knows that X is the thing for her to do.

Not only is (LA) implausible, things look bad for The Knowledge Rule if it has to rely on (LA) being true. None of the defenders of The Knowledge Rule has given us an argument that (LA) is true. One of them has given us all we need to show that (LA) is false! It doesn't look like the kind of principle that The Knowledge Rule should have to depend upon. So, defending The Knowledge Rule here looks hopeless.

Note that given premise 1 of the Master Argument, as corrected, *every* instance of (LA) has to be true for The Knowledge Rule to be universally true. Let's say that you thought (LA) was true when X is *starting a war*, but not when X is *buying flood insurance*. Then we can use the case of Raj and Nik to show that The Knowledge Rule fails, since Raj can say that buying flood insurance is what to do in a case where it is what to do, but he doesn't know this.

One final observation about the anti-luminosity argument. Given the way Williamson presents the anti-luminosity argument, it can appear that in all but a few cases, if p , the salient agent can know that p . After all, the only examples Williamson gives are cases that are only picked out by something like the Least Number Theorem. So, one might think that while luminosity principles are false, they are approximately true. More precisely, one might think that in all but a few weird cases near the borderline, if p , then a salient agent is in a position to know p . If so, then the failures of luminosity aren't of much practical interest, and hence the failures of The Knowledge Rule we've pointed out aren't of much practical interest.

We think this is all mistaken. Luminosity failures arise because agents have less than infinite discriminatory capacities. The worse the discriminatory capacities, the greater the scope for luminosity failures. When agents have very poor discriminatory capacities, there will be very many luminosity failures. This is especially marked in decision-making concerning war. The fog of war is thick. There is very much that we don't know, and what we do know is based on evidence that is murky and ephemeral. There is very little empirical information that we know that we know.

If there are certain actions (such as starting a war) that are proper only if we know a lot of empirical information, the general case will be that we cannot know that these actions are correct, even when they are. This suggests that luminosity failures, where an action is correct but not known to be correct, or a fact is known but not known to be known, are not philosophical curiosities. In epistemically challenging environments, like a war zone, they are everyday facts of life.

4 Moore's Paradox

There is a standard argument for The Knowledge Rule that goes as follows. First, if the Knowledge Rule did not hold, then certain Moore paradoxical assertions would be acceptable. In particular, it would be acceptable to assert *q*, *but I don't know that q*.⁴ But second, Moore paradoxical assertions are never acceptable. Hence, The Knowledge Rule holds. We reject both premises of this argument.

To reject the first premise, it suffices to show that some rule other than The Knowledge Rule can explain the unacceptability of Moore paradoxical assertions. Consider, for example, The Undefeated Reason rule.

The Undefeated Reason Rule Assert that *p* only if you have an undefeated reason to believe that *p*.

The Undefeated Reason Rule says that *q but I don't know that q* can be asserted only if the speaker has an undefeated reason to believe it. That means the speaker has an undefeated reason to believe each conjunct. That means that the speaker has an undefeated reason to believe that they don't know *q*. But in every case where it is unacceptable to both assert *q* and assert that you don't know *q*, the speaker's undefeated reason to believe they don't know *q* will be a defeater for her belief that *q*. If you have that much evidence that you don't know *q*, that will in general defeat whatever reason you have to believe *q*.

We don't claim that The Undefeated Reason Rule is correct. (In fact, we prefer the rules we'll discuss in section 5.) We do claim that it provides an alternative explanation of the unacceptability of instances of *q but I don't know that q*. So, we claim that it undermines the first premise of Williamson's argument from that unacceptability to The Knowledge Rule.

We also think that Williamson's explanation of Moore paradoxicality over-generates. There is generally something odd about saying *q but I don't know that q*. We suspect that the best explanation for why this is odd will be part of a broader explanation that also explains, for instance, why saying *I promise to do X, but I'm not actually doing to do X* is also defective. Williamson's explanation isn't of this general form. He argues that saying *q but I don't know that q* is defective because it is defective in every context to both assert *q* and assert that you don't know that *q*. But we don't think that it is always defective to make both of these assertions.⁵ In particular, if a

⁴(Williamson, 2000, §11.3), for instance, shows the strength of this argument.

⁵This is why we hedged a little two paragraphs ago about what precisely The Undefeated Reason Rule explains. We suspect that many in the literature have misidentified the explicandum.

speaker is asked whether q is true, and whether they know that q , it can be acceptable to reply affirmatively to the first question, but negatively to the second one. If so, then the second premise of Williamson's argument from Moore paradoxicality is also false.

Imagine that the Indian Prime Minister is a philosopher in her spare time. After the big speech to Parliament she goes to her Peninsula Reading Group. It turns out Michael Walzer and Tim Williamson are there, and have questions about the speech.

TW: Do you agree that knowledge requires safety?
 PM: Yes, yes I do.
 TW: And do you agree that your belief that going to war is the thing to do is not safe?
 PM: Right again.
 TW: So, you don't know that going to war is the thing to do?
 PM: You're right, I don't.
 MW: But is it the thing to do?
 PM: Yes.

The Prime Minister's answers in this dialogue seem non-defective to us. But if Williamson's explanation of why Moore paradoxical utterances are defective is correct, her answers should seem defective. So, Williamson's explanation over-generates. Whether or not it is true that all assertions of sentences of the form q *but I don't know that q* are defective, it isn't true that there is a defect in any performance that includes both an assertion of q and an assertion of the speaker's ignorance as to whether q . The Prime Minister's performance in her reading group is one such performance. So, the explanation of Moore paradoxicality cannot be that any such performance would violate a norm governing assertion.

To sum up, then, we've argued that The Knowledge Rule (a) fails to be the only explanation of Moore paradoxicality, and (b) misclassifies certain performances that are a little more complex than simple conjunctive assertions as defective. So, there's no good argument from Moore paradoxicality to The Knowledge Rule.

5 Action and Assertion

If we're right, there's a striking asymmetry between certain kinds of assertions. In the war example, early in her speech, the Prime Minister says (2).

(2) The enemy has been murdering his own civilians.

That's not the kind of thing she could properly say if it could easily have been false given her evidence. And like many assertions, this is not an assertion whose appropriateness is guaranteed by its truth. Asserting (2) accuses someone of murder, and you can't properly make such accusations without compelling reasons, even if they happen to be true. On the other hand, we say, the truth of (1) does (at least when it is accepted on the right basis) suffice to make it properly assertable.

There's a similar asymmetry in the flood insurance example. In that example, (3) is true, but neither Raj nor Nik knows it.

(3) Raj and Nik's business will not flood this year.

Again, in these circumstances, this isn't the kind of thing Raj can properly say. Even though (3) is true, it would be foolhardy for Raj to make such a claim without very good reasons. By contrast, again, we say that Raj can properly assert that the thing to do, in their circumstances, is to buy flood insurance, even though he does not know this.

There are two directions one could go at this point. If we're right, any proposed theory of the norms governing assertion must explain the asymmetry. Theories that cannot explain it, like The Knowledge Rule, or the Certainty Rule proposed by Jason Stanley (2008), or the Rational Credibility Rule proposed by Igor Douven (2006), are thereby refuted.

The Certainty Rule Assert only what is certain.

The Rational Credibility Rule Assert only what is rationally credible.

The Certainty Rule fails since the Prime Minister is not certain of (1). And the Prime Minister can't be certain of (1), since certainty requires safety just as much as knowledge does.

It's a little harder to show our example refutes The Rational Credibility Rule. Unlike knowledge, a safety constraint is not built into the concept of rational credibility. (Since rational credibility does not entail truth, in Douven's theory, it can hardly entail truth in nearby worlds.) But we think that safety constraints may still apply to rational credibility in some particular cases. If you aren't very good at judging building heights of tall buildings to a finer grain than 10 meters, then merely looking at a building that is 84 meters tall does not make it rationally credible for you that the building is more than 80 meters tall. In general, if your evidence does not give you much reason to think you are not in some particular world where p is false, and you didn't have prior reason to rule that world out, then p isn't rationally credible. So, when evidence doesn't discriminate between nearby possibilities, and p is false in nearby possibilities, p isn't rationally credible.

And that, we think, is what happens in our two examples. Just as someone looking at an 84 meter building can't rationally credit that it is more than 80 meters tall, unless they are abnormally good at judging heights, agents for whom X is just barely the thing to do can't rationally credit that X is the thing to do. By The Rational Credibility Rule, they can't say X is the thing to do. But they can say that; that's what our examples show. So, The Rational Credibility Rule must be wrong.

But we can imagine someone pushing in the other direction, perhaps with the help of this abductive argument.

1. A speaker can only assert things like (2) or (3) if they know them to be true.
2. The best explanation of premise 1 of this argument is The Knowledge Rule.
3. So, The Knowledge Rule is correct.

This isn't a crazy argument. Indeed, it seems to us that it is implicit in some of the better arguments for The Knowledge Rule. But we think it fails. And it fails because there are alternative explanations of the first premise, explanations that don't make mistaken predictions about the Prime Minister's speech. For instance, we might have some kind of Evidence Responsiveness Rule.

The Evidence Responsiveness Rule Assert that p only if your attitude towards p is properly responsive to the evidence you have that bears on p .

Given how much can be covered by 'properly', this is more of a schema than a rule. Indeed, it is a schema that has The Knowledge Rule as one of its precisifications. In *Knowledge and Its Limits*, Williamson first argues that assertion is "governed by a non-derivative evidential rule" (249), and then goes on to argue that the proper form of that rule is The Knowledge Rule. We agree with the first argument, and disagree with the second one.⁶

Note that even a fairly weak version of The Evidence Responsiveness Rule would explain what is going on with cases like (1) and (2). Starting a war is a serious business. You can't properly do it *unless* your views about the war are evidence responsive in the right way. You can't, that is, correctly *guess* that starting the war is the thing to do. You *can* correctly guess that starting the war will be utility maximizing. And you can correctly guess that starting the war would be what to choose if you reflected properly on the evidence you have, and the moral significance of the choices in front of you. But you simply can't guess that starting the war is what to do, and be right. If you're merely guessing that starting a war is thing to do, then you're wrong to start that war. So, if (1) is true, and the Prime Minister believes it, her belief simply *must* be evidence responsive. Then, by The Evidence Responsiveness Rule, she can assert it.

For most assertions, however, this isn't the case. Even if it's true that it will rain tomorrow, the Prime Minister's could believe that without her belief being evidence responsive. In general, p does not entail that S even believes that p , let alone that this belief of S 's is evidence responsive. But in cases like (1), this entailment does hold, and that's what explains the apparent asymmetry that we started this section with.

The Evidence Responsiveness Rule also handles so called 'lottery propositions' nicely. If you know that the objective chance of p being true is c , where c is less than 1, it will seem odd in a lot of contexts to simply assert p . In his arguments for The Knowledge Rule, Williamson makes a lot of this fact. In particular, he claims that the best explanation for this is that we can't know that p on purely probabilistic grounds. This has proven to be one of the most influential arguments for The Knowledge Rule in the literature. But some kind of Evidence Responsiveness Rule seems to handle

⁶Actually, our agreement with Williamson here is a bit more extensive than the text suggests. Williamson holds that part of what makes a speech act an *assertion* as opposed to some other kind of act is that it is governed by The Knowledge Rule. Although many philosophers agree with Williamson that The Knowledge Rule is true, this fascinating claim about the metaphysics of speech acts has been largely ignored. Translating Williamson's work into the terminology of this paper, we're inclined to agree that a speech act is an assertion partly in virtue of being responsive to evidence in the right way. But filling in the details on this part of the story would take us too far from the main storyline of this paper.

lottery cases even more smoothly. In particular, an Evidence Responsiveness Rule that allows for what constitutes ‘proper’ responsiveness to be sensitive to the interests of the conversational participants will explain some odd features concerning lottery propositions and assertability.

In the kind of cases that motivate Williamson, we can’t say p where it is objectively chancy whether p , and the chance of p is less than 1. But there’s one good sense in which such an assertion would not be properly responsive to the evidence. After all, in such a case there’s a nearby world, with all the same laws, and with all the same past facts, and in which the agent has all the same evidence, in which p is false. And the agent knows all this. That doesn’t look like the agent is being properly responsive to her evidence.

On the other hand, we might suspect that Williamson’s arguments concerning lottery propositions overstate the data. Consider this old story from David Lewis (1996).⁷

Pity poor Bill! He squanders all his spare cash on the pokies, the races, and the lottery. He will be a wage slave all his days . . . he will never be rich. (Lewis, 1996, 443 in reprint)

These seem like fine assertions. One explanation of the appropriateness of those assertions combines The Knowledge Rule with contextualism about assertion.⁸ But contextualism has many weaknesses, as shown in Hawthorne (2004) and Stanley (2005). A less philosophically loaded explanation of Lewis’s example is that proper responsiveness comes in degrees, and for purposes of talking about Bill, knowing that it’s overwhelmingly likely that he’s doomed to wage slavery is evidence enough to assert that he’ll never be rich. The details of this explanation obviously need to be filled in, but putting some of the sensitivity to conversational standards, or practical interests, into the norms of assertion seems to be a simpler explanation of the data than a contextualist explanation. (It would be *a priori* quite surprising if the norms of proper assertion were not context-sensitive, or interests-sensitive. The norms of appropriateness for most actions are sensitive to context and interests.) So The Evidence Responsiveness Rule seems more promising here than The Knowledge Rule.

A harder kind of case for The Knowledge Rule concerns what we might call ‘academic assertions’. This kind of case is discussed in Douven (2006) and in Maitra (2010). In academic papers, we typically make assertions that we do not know. We don’t know that most of the things we’ve said here are true. (Before the last sentence we’re not sure we knew that any of the things we said were true.) But that’s because knowledge is a bad standard for academic discourse. Debate and discussion would atrophy if we had to wait until we had knowledge before we could present a view. So, it seems that assertion can properly outrun knowledge in academic debate.

⁷We’ve slightly modified the case. Lewis says we can say that we *know* Bill will never be rich. That seems to us to be a much more controversial than what we’ve included here.

⁸The combination is slightly trickier to state than would be ideal. The explanation we have in mind is that S can properly assert p only if S can truly say *I know that p* , where ‘know’ in this utterance is context sensitive.

Again, a context-sensitive version of The Evidence Responsiveness Rule explains the data well. Although you don't need to *know* things to assert them in philosophy papers, you have to have evidence for them. We couldn't have just spent this paper insisting louder and louder that The Knowledge Rule is false. We needed to provide evidence, and hopefully we've provided a lot of it. In some contexts, such as testifying in court, you probably need more evidence than what we've offered to ground assertions. But in dynamic contexts of inquiry, where atrophy is to be feared more than temporary mistakes, the standards are lower. Good evidence, even if not evidence beyond any reasonable doubt, or even if not enough for knowledge, suffices for assertion. That's the standard we typically hold academic papers to. Like with lotteries, we think the prospects of explaining these apparently variable standards in terms of a norm of assertion that is context-sensitive are greater than the prospects for explaining them in terms of contextually sensitive knowledge ascriptions.

Here's a different and somewhat more speculative proposal idea for a rule that also explains the asymmetry we started this section with. We call it the Action Rule.

The Action Rule Assert that p only if acting as if p is true is the thing for you to do.

We take the notion of acting as if something is true from Stalnaker (1973). Intuitively, to act as if p is true is to build p into one's plans, or to take p for granted when acting. This, note, is not the same as using p as a basis for action. When Raj buys flood insurance, he acts as if buying flood insurance is the thing to do. But the fact that buying flood insurance is the thing to do isn't the basis for his action. (Since he does not know this, one might suspect it wouldn't be a good basis.) Instead his basis is what he knows about the river, and his business, and its vulnerability to flooding. When an agent is trying to maximise the expected value of some variable (e.g., utility, profit, etc.), then to act as if p is true is simply to maximise the conditional expected value of that variable, in particular, to maximise the expected value of that variable conditional on p . Even when one is not maximising any expected value, we can still use the same idea. To act as if p is to take certain conditional obligations or permissions you have – in particular, those obligations or permissions that are conditional on p – to be actual obligations or permissions.

To see how The Action Rule generates the intended asymmetry, we'll need a bit of formalism. Here are the terms that we will use.

- X denotes an action, agent, circumstance triple $\langle X_{Action}, X_{Agent}, X_{Circumstance} \rangle$. We take such triples to have a truth value. X is true iff X_{Agent} performs X_{Action} in $X_{Circumstance}$.
- ThingToDo(X) means that X is the thing to do for X_{Agent} in $X_{Circumstance}$.
- Act(S, p) means that agent S acts as if p is true.
- Assert(S, p) means that agent S can properly assert that p .

So, The Action Rule is this.

$$\text{Assert}(S, p) \rightarrow \text{ThingToDo}(\text{Act}(S, p))$$

In our derivations, the following equivalence will be crucial.

$$\text{Act}(X_{Agent}, \text{ThingToDo}(X)) \leftrightarrow X$$

That is, acting as if X is what to do (in your circumstances) is simply to do X (in those circumstances). And in doing X , you're acting as if X is what to do (in your circumstances). We take this equivalence to be quite resilient; in particular, it holds under operators like 'ThingToDo'. So, adding that operator to the previous equivalence, we get another equivalence.

$$\text{ThingToDo}(\text{Act}(X_{Agent}, \text{ThingToDo}(X))) \leftrightarrow \text{ThingToDo}(X)$$

If we substitute $\text{ThingToDo}(X)$ for p in The Action Rule, we get this.

$$\text{Assert}(X_{Agent}, \text{ThingToDo}(X)) \rightarrow \text{ThingToDo}(\text{Act}(X_{Agent}, \text{ThingToDo}(X)))$$

But by the equivalence we derived earlier, that's equivalent to the following.

$$\text{Assert}(X_{Agent}, \text{ThingToDo}(X)) \rightarrow \text{ThingToDo}(X)$$

So, we get the nice result that The Action Rule is trivially satisfied for any true claim about what is to be done. That is, for the special case where p is *X is the thing for you to do*, The Action Rule just reduces to something like the Truth Rule. And so we get a nice explanation of why the Prime Minister and Raj can properly make their assertions about what to do in their respective circumstances.⁹

To explain the other side of the asymmetry with which we began this section, note that these biconditionals do not hold where p is an arbitrary proposition, and S an arbitrary agent.

$$\begin{aligned} \text{ThingToDo}(\text{Act}(S,p)) &\leftrightarrow p \\ \text{Act}(S, \text{ThingToDo}(\text{Act}(S,p))) &\leftrightarrow p \end{aligned}$$

To see this, let p be the proposition expressed by (4). To act as if this is true is to, *inter alia*, not buy flood insurance. If there won't be a flood, buying flood insurance is throwing away money, and when you're running a business, throwing away money isn't the thing to do. In symbols, $\text{Act}(\text{Raj and Nik}, p)$ is equivalent to *Raj and Nik don't buy flood insurance*. But not buying flood insurance is not the thing to do. The prudent plan is to buy flood insurance. So, $\text{ThingToDo}(\text{Act}(\text{Raj and Nik}, p))$ is false, even though p is true. So, the first biconditional fails. Since Raj and Nik do go on to buy flood insurance, i.e., since they don't act as if $\text{ThingToDo}(\text{Act}(\text{Raj and Nik}, p))$, the left-hand-side of the second biconditional is also false. But again, the right-hand-side is true. So, that biconditional is false as well. And without those biconditionals, The Action Rule doesn't collapse into $\text{Assert}(S,p) \rightarrow p$.

⁹The derivation here is deliberately simplified in one way. We haven't included anything about the *bases* for action or assertion. We don't think being sensitive to bases in the formalism would make a material change, but it would obscure the structure of the argument.

We have thus far argued that The Action Rule can provide an explanation for the asymmetry we noted at the beginning of this section.¹⁰ This is not, however, meant to be anything like a complete defence of that rule. That would require a lot more than we've provided here. But we do think that the Action Rule can explain a lot of the phenomena that are meant to motivate The Knowledge Rule, as well as some phenomena The Knowledge Rule struggles with. But we do think The Action Rule has some virtues. We'll close with a discussion of how it explains the two kinds of cases that we argued that The Evidence Responsiveness Rule handles well.

To see this, consider first 'lottery propositions'. If you know that the objective chance of p being true is c , where c is less than 1, it will seem odd in a lot of contexts to simply assert p . In his arguments for The Knowledge Rule, Williamson makes a lot of this fact. In particular, he claims that the best explanation for this is that we can't know that p on purely probabilistic grounds. This has proven to be one of the most influential arguments for The Knowledge Rule in the literature.

We suggest that The Action Rule can offer an alternative a nice explanation for why it's often defective to assert lottery propositions. Note first that in a lot of cases, it isn't rational for us to act on p when we have only purely probabilistic evidence for it, especially when acting on p amounts to betting on p at sufficiently unfavourable odds. This point is something of a staple of the 'interest-relative-invariantism' literature on knowledge.¹¹ To take a mundane case, imagine that you're cleaning up your desk, and you come across some lottery tickets. Most are for lotteries that have passed, that you know you lost. One ticket, however, is for a future lottery, which you know you have very little chance of winning. In such a case, to act as if the ticket for the future lottery would lose would be to throw it out along with the other tickets. But that would be irrational, and not at all how we'd act in such a case. That is to say, in such a case, we don't (and shouldn't, rationally speaking) act as if the ticket for the future lottery will lose, even though we take that outcome to be highly probable.

If acting as if a lottery proposition is true isn't the thing to do, then The Action Rule will say that asserting such a proposition is defective. Therefore, we think that The Action Rule can capture why in many cases you can't in general assert lottery propositions.

A harder kind of case for The Knowledge Rule concerns what we might call 'academic assertions'. This kind of case is discussed in Douven (2006) and in Maitra

¹⁰This explanation makes some interestingly different predictions from the explanation in terms of The Evidence Responsiveness Rule. Suppose that for relatively trivial decisions, like where to go for a walk on a nice summer day, one can correctly guess that X is the thing to do. Then the Evidence Responsiveness Rule would suggest that the truth of claims about where to go for a walk is not sufficient grounds for their assertability, while the Action Rule would still imply that truth is sufficient grounds for assertability.

We're not sure that this supposition – that for relatively trivial decisions, one can correctly guess that X is the thing to do – is coherent, nor what to say about assertability judgments in (imagined) cases where the supposition holds. So, we're not sure we can really use this to discriminate between the two proposed explanations. Nevertheless, it is interesting to note how the explanations come apart. Thanks here to Susanna Schellenberg.

¹¹See, for instance, Fantl and McGrath (2002), Hawthorne (2004), Stanley (2005), and Weatherson (2005).

(2010). In academic papers, we typically make assertions that we do not know. We don't know that most of the things we've said here are true. (Before the last sentence we're not sure we knew that any of the things we said were true.) But that's because knowledge is a bad standard for academic discourse. Debate and discussion would atrophy if we had to wait until we had knowledge before we could present a view. So, it seems that assertion can properly outrun knowledge in academic debate.

Academic assertions raised a problem for The Knowledge Rule because proper assertion in the context of inquiry can outrun knowledge. But note that action in such a context can also properly outrun knowledge. It would slow down learning dramatically if people didn't engage in various projects that really only make sense if some hypothesis is true. So, academics will study in archives, conduct experiments, write papers, etc. etc., and do so on the basis of reasons they no more know than we know the truth of the speculative claims of this paper. And this is all to the good; the alternative is a vastly inferior alternative to academia as we know it. So, in some fields, action requires much less than knowledge. Happily, in those fields, assertion also requires much less than knowledge. Indeed, the shortfalls in the two cases seem to parallel nicely. And this parallel is neatly captured by The Action Rule.

As we said, none of this is a knockdown case for The Action Rule. Our primary purpose is to argue against The Knowledge Rule. As long as the Action Rule is plausible, we have defeated the abductive argument for The Knowledge Rule that was discussed at the start of this section, and we think we've done enough to show it is plausible. We also hope we've made a successful case for moving the study of assertability away from rules like The Knowledge Rule, and instead have it be more tightly integrated with our best theories about evidence and action.

References

- Bach, Kent. 2010. "Knowledge In and Out of Context." In Joseph Keim Campbell, Michael O'Rourke, and Harry S. Silverstein (eds.), *Knowledge and Skepticism*, unknown. Cambridge, MA: MIT Press.
- Black, Max. 1952. "Saying and Disbelieving." *Analysis* 13:25–33.
- Cappelen, Herman and Lepore, Ernest. 2005. *Insensitive Semantics: A Defence of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell.
- DeRose, Keith. 2002. "Assertion, Knowledge and Context." *Philosophical Review* 111.
- Douven, Igor. 2006. "Assertion, Knowledge and Rational Credibility." *Philosophical Review* 115.
- Fantl, Jeremy and McGrath, Matthew. 2002. "Evidence, Pragmatics, and Justification." *Philosophical Review* 111:67–94.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hindriks, Frank. 2007. "The Status of the Knowledge Account of Assertion." *Linguistics and Philosophy* 30:393–406.

- Lewis, David. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74:549–567. Reprinted in *Papers in Metaphysics and Epistemology*, pp. 418–446.
- Maitra, Ishani. 2010. "Assertion, Norms and Games."
- Sainsbury, Mark. 1996. "Vagueness, Ignorance and Margin for Error." *British Journal for the Philosophy of Science* 46:589–601.
- Stalnaker, Robert. 1973. "Presuppositions." *Journal of Philosophical Logic* 2.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford University Press.
- . 2008. "Knowledge and Certainty." *Philosophical Issues* 18:35–57.
- Weatherson, Brian. 2004. "Luminous Margins." *Australasian Journal of Philosophy* 82:373 – 383.
- . 2005. "Can We Do Without Pragmatic Encroachment?" *Philosophical Perspectives* 19:417–443.
- Weiner, Matthew. 2005. "Must We Know What We Say." *Philosophical Review* 114:227–251.
- Williamson, Timothy. 1996. "Knowing and Asserting." *Philosophical Review* 105:489–523.
- . 2000. *Knowledge and its Limits*. Oxford University Press.