

Assertion, Knowledge, and Action

Ishani Maitra and Brian Weatherson

February 28, 2009

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 else 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.

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 that one of the more attractive alternatives is The Action Rule.

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

¹ Timothy Williamson (2000: Chapter 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 (2007) 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 you justifiably believe. Kent Bach (forthcoming) 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, as Williamson thinks, or in virtue of other norms, 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.

We're not going to suggest that there is a conclusive argument for The Action Rule. But we will make three claims about this rule. First, it explains an interesting asymmetry that other rules cannot explain. Second, it undermines another positive argument for The Knowledge Rule, namely the argument from lottery propositions. Finally, it does not collapse into other rules, such as The Knowledge Rule or a rule about justified belief.

1. Speaking about What to Do

We start by motivating premise 1 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.

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

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. 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). Note that she doesn't violate *any* norms when she actually launches the war. By hypothesis, starting the war is the thing for her to do, all things considered. So it is proper that she starts it. Her only norm violation is when she (truly) asserts that going to war is what to do.

Now it would be absurd to think that a Prime Minister could launch a war without making any speech about this to the House. So what kind of speech would be compatible with Williamson's rules? The best she could do would be to recite the evidence supporting the war, and then end with something like (2).

- (2) These are all strong reasons for going to war, but I can't say whether they make going to war the thing to do. Nevertheless, I'm starting the war tomorrow.

We think that (1), which violates The Knowledge Rule, is a much more appropriate end to the speech than (2), which does not violate it. And (2) is a better way to end the speech than anything else we can think of that does not violate The Knowledge Rule.

To put the point another way, the Prime Minister has three options regarding how to end her speech. One option is to lay out the reasons for war, and then end without addressing the issue of whether going to war is the thing to do. Such a speech would satisfy The Knowledge Rule, but put the Prime Minister in the position of not being able to say what her reasons are reasons for. A second option is to end with

something like (2), which does address the issue, but is pragmatically unhappy, to put it mildly. The third, and – we think – clearly best option is to end the speech with something like (1). In so doing, the Prime Minister connects her reasons with the action they’re reasons for, and in so doing, offers the complete justification for her decision to go to war. Unfortunately, this option is ruled out by The Knowledge Rule. More generally, it seems to us that once the Prime Minister has looked carefully enough at the evidence, and thought carefully enough about what’s at issue, she has done enough to be able to properly say that starting the war is what to do in her circumstances.

This first example is one in which moral considerations are, if not the only considerations at issue, at least very important in coming to an all-things-considered judgment about what to do. But, as we’ll see next, the same pattern comes up in prudential cases as well.

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 a pragmatic tension, to put it mildly, in Raj’s actions. 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. So The Knowledge Rule is wrong.

2. Bases for Action and Assertion

One might worry that premise 1 in our master argument is mistaken. 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 that it is 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 right thing for Raj* to do. On the 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 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 what Raj* does in this variant example is wrong, 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 Chapter 4 of Williamson (2000).) 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 is 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. So 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 – Indalia’s leaders, Raj and Nik – 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 $\sim p$, they can’t know p . Using variants of the Indalia case, 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, in our Indalia example, a crucial piece of the 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 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 supplying misinformation. 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. So if the spy tells them that p , they come to know that p , and they can discriminate p worlds from $\sim p$ 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 , they don’t have sufficient reason to go to war. So, given that they don’t know the spy is reliable, they don’t know that they have sufficient reason to go to war. But the spy really is reliable, so they do know p , so they do have sufficient reason to go to war.

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 we do a sample, and we conclude that the enemy has killed n of his citizens, with a margin of error of m . So, assuming there really are n killings, we know that he’s killed between $n - m$ and $n + m$ of his citizens. Since knowing that he’s killed $n - m$ people is sufficient reason to start a war, the war can be properly started. But all this is consistent with $n - 2m$ deaths not being sufficient reason to start a war, especially if m is large. Since for all we know, the enemy has killed only $n - m$ people, and we can’t tell how many people he’s killed to a closer approximation than m , for all we know we know, he’s killed only $n - 2m$ people. So for all we know we know, we don’t have sufficient reason to start a war. But that’s consistent with actually having sufficient reason to start a war.

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. So 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 needs. 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. That means there is much we know that is based on evidence we don't know we have. But that means there is much we know that we don't know we know. Given how much can go wrong, epistemically speaking, in wartime, we end up knowing very little. That suggests our margins of error are high in wartime. And working through the models Williamson provides shows that when that is true, there is a lot that we know without knowing that we know. So we think these luminosity failures are not philosophical curiosities; in complicated practical settings, they may be closer to the norm.

4. Moore's Paradox

There is a standard argument for The Knowledge Rule, namely, that if it 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* .³ But 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 Undeclared Reason rule.

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

The Undeclared 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 that, we think, *is* a defeater for her belief that q . If you have that much evidence that you don't know q , that defeats whatever reason you have to believe q .

We don't claim that The Undeclared Reason Rule is correct. We do claim that it provides an alternative explanation of the unacceptability of q , *but I don't know that q* . So we claim that it undermines any argument from that unacceptability to The Knowledge Rule.

We also reject the general premise that Moore paradoxical assertions are always defective. Imagine the Prime Minister in the war example being interviewed on *Newsnight*.

Interviewer: Is launching this war what we must do?

³ Williamson (2000: Section 11.3), for instance, shows the strength of this argument.

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 thing to do?

Prime Minister: Well, you almost never *know* that launching a war is the thing to do. 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 thing to do, and when we look back at this war with the benefit of hindsight, fair critics 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 going to war is what to do. Maybe we are just being too generous because of this graciousness, but it seems to us that the Prime Minister's little speech is not thereby defective in the way that Moore paradoxical utterances are alleged to be always defective.

So we have a fairly flat-footed answer to the argument for a knowledge norm from facts about Moore paradoxicality. Whether or not it is true that all sentences of the form *q but I don't know that q* are defective, it isn't true that all speeches that assert that *q*, while also asserting that *q* is not known are defective. The Prime Minister's answer is one such speech. So the explanation of Moore paradoxicality cannot be that uttering any such claim would violate a norm governing assertion.

5. Action and Assertion

There are some odd asymmetries in the two cases we focused on earlier. In the war example, early in her speech, the Prime Minister says (3).

(3) 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 (3) 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, (4) is true, but neither Raj nor Nik knows it.

(4) 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 (4) 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.

In both cases, then, there are asymmetries between what it takes to properly assert claims of the form *X is the thing to do*, and what it takes to properly make other sorts of assertions. The Knowledge Rule cannot explain these asymmetries. What kind of theory could explain it?

Clearly a rule stronger than The Knowledge Rule cannot explain it. Jason Stanley (2007), for instance, endorses The Certainty Rule.

The Certainty Rule: Assert only what is certain.

But since the Prime Minister is not certain of (1), that's not plausible. And the Prime Minister can't be certain of (1), since certainty requires safety just as much as knowledge does. The same argument tells against some weaker rules. For instance, Igor Douven (2006) endorses The Rational Credibility Rule.

The Rational Credibility Rule: Assert only what is rationally credible.

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 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. So 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.

What we need is a rule that discriminates between claims of the form *X is the thing to do*, and more prosaic claims about everyday matters. We suggest that The Action Rule satisfies this constraint.

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

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.⁴ 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.

To explain the other side of the asymmetries 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 $\text{Assert}(S,p) \rightarrow p$.

We have thus far argued that The Action Rule can provide an explanation for the asymmetries we noted at the beginning of this section. This is not, however, meant to be anything like a complete defence of that rule. In particular, we haven't shown that there aren't other rules that can explain these asymmetries as well. Nor have we shown that The Knowledge Rule doesn't govern *some* assertions, for it may very well be (and we think it's plausible) that for certain values of p , acting as if p is the thing for someone to do only if they know that p . What we deny, however, is that that is true for all values of p . In particular, as we've argued, it can't be true when p takes the value *X is the thing to do*, unless *being the thing to do* is luminous.

⁴ If the agent correctly believes that X is the thing to do, but believes this for the wrong reasons, then they won't have good reason to believe that The Action Rule is satisfied, so arguably they shouldn't say that X is the thing to do.

In the remainder of the paper, we will offer some further clarifying remarks about The Action Rule, and then close by mentioning what we take to be some further virtues of that rule.

The Action Rule would collapse into some familiar rules we have already dismissed if we posit some extra rules about action. Consider the following Knowledge-Action Principle (KAP).

(KAP) Acting as if p is true is the thing to do only if the agent knows that p

A similar, though we think crucially different, principle is defended in Hawthorne and Stanley (2008). We can see that (KAP) is false by substituting *X is the thing to do* for p . We argued above that when we make that substitution, the left-hand-side of (KAP) reduces to *X is the thing to do*. But the right-hand-side of (KAP) is *The agent knows that X is the thing to do*. So under this substitution, (KAP) becomes a luminosity principle for *being the thing to do*. But *being the thing to do* is not luminous. So (KAP) is false. (We leave as a question for future investigation whether this kind of argument can be generalised into an argument against the principle Hawthorne and Stanley do defend.)

Alternatively, consider the following Functionalist Theory of Belief (FTB).

(FTB) An agent believes that p iff they act as if p is true.

Arguably it follows from (FTB) that an agent rationally believes p iff they rationally act as if p is true. Assume, as seems plausible, that it is rational to act as if p is true iff acting as if p is true is the thing to do. Then (FTB) gives us that acting as if p is true is the thing to do iff the agent rationally believes that p . And that means that The Action Rule is equivalent to The Rational Credibility Rule. That would be bad, since we argued above that The Rational Credibility Rule is false.

Fortunately, the problem lies with (FTB), not with The Action Rule. Again, substitute *X is the thing to do* for p . We get that the agent believes that *X is the thing to do* iff they act as if *X is the thing to do*. So they rationally believe *X is the thing to do* iff they rationally do *X*. But that's false in cases like Raj buying flood insurance. It's a borderline case, so he can't rationally believe that buying flood insurance is the thing to do. But it's a borderline case where buying flood insurance is the thing to do (by hypothesis), so he can rationally buy flood insurance. So (FTB) is false, and there is no danger of The Action Rule collapsing into The Rational Credibility Rule.

We'll end with discussion of two kinds of cases that have played a prominent role in recent debates. We won't say much about the two cases, because we've discussed them at some length elsewhere – though not in the context of The Action Rule.⁶

The first kind of case concerns 'lottery propositions'. If you know the objective chance of p being true is c , where c is less than 1, it seems odd in a lot of contexts to simply assert p . In his arguments for The Knowledge Rule, Williamson makes a lot of this fact, claiming that the best explanation for this is that we can't know that p on purely probabilistic grounds. But it's also true in a lot of cases that we can't act on p when we have 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. See, for instance, Fantl and McGrath (2002), Hawthorne (2004), Stanley (2005), and Weatherson (2005). So we think that The Action Rule is consistent with the fact that you can't in general assert lottery propositions.

The second kind of case concerns what we might call 'academic assertions'. This kind of case is discussed in Douven (2006) and in (Author Paper). 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.

But note that action 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. But we do think there is a case for moving the study of assertability away from rules like The Knowledge Rule, and instead have it be more tightly integrated with general principles about action.

References

- Bach, Kent (forthcoming) "Knowledge in and out of Context" in *Knowledge and Skepticism*, Joseph Keim Campbell and Michael O'Rourke (eds.), Cambridge (MA): MIT Press.
- Black, Max (1952) "Saying and Disbelieving" *Analysis* 13: 25-33.
- DeRose, Keith (2002) "Assertion, Knowledge, and Context" *Philosophical Review* 111: 167-203.
- Douven, Igor (2006) "Assertion, Knowledge, and Rational Credibility" *Philosophical Review* 115: 449-85.
- Fantl, Jeremy and Matthew McGrath (2002) "Evidence, Pragmatics and Justification" *Philosophical Review* 111: 67-94.
- Gibbard, Allan (2003) *Thinking How to Live*, Cambridge (MA): Harvard University Press.
- Hawthorne, John (2004) *Knowledge and Lotteries*, Oxford: Oxford University Press.

⁶ References suppressed for blind refereeing.

- Hawthorne, John and Jason Stanley (2008) "Knowledge and Action" *Journal of Philosophy* 105: 571-90.
- Hindriks, Frank (2007) "The Status of the Knowledge Account of Assertion" *Linguistics and Philosophy*, 30: 393-406.
- Sainsbury, Mark (1996) "Vagueness, Ignorance and Margin for Error" *British Journal for the Philosophy of Science*, 46: 589-601.
- Stanley, Jason (2005) *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Stanley, Jason (2007) "Knowledge and Certainty" *Philosophical Issues* 18: 33-55.
- Weatherson, Brian (2004) "Luminous Margins" *Australasian Journal of Philosophy*, 83: 373-83.
- Weiner, Matthew (2005) "Must We Know What We Say" *Philosophical Review*, 114: 227-51.
- Williamson, Timothy (1996) "Knowing and Asserting" *The Philosophical Review* 105: 489-523.
- Williamson, Timothy (2000) *Knowledge and Its Limits*, Oxford: Oxford University Press.