

Scepticism, Rationalism and Externalism^{*}

This paper is about three of the most prominent debates in modern epistemology. The conclusion is that three *prima facie* appealing positions in these debates cannot be held simultaneously.

The first debate is **scepticism vs anti-scepticism**. My conclusions apply to *most* kinds of debates between sceptics and their opponents, but I will focus on the inductive sceptic, who claims we cannot come to know what will happen in the future by induction. This is a fairly weak kind of scepticism, and I suspect many philosophers who are generally anti-sceptical are attracted by this kind of scepticism. Still, even this kind of scepticism is quite unintuitive. I'm pretty sure I know (1) on the basis of induction.

(1) It will snow in Ithaca next winter.

Although I am taking a very strong version of anti-scepticism to be intuitively true here, the points I make will generalise to most other versions of scepticism. (Focussing on the inductive sceptic avoids some potential complications that I will note as they arise.)

The second debate is a version of **rationalism vs empiricism**. The kind of rationalist I have in mind accepts that some deeply contingent propositions can be known a priori, and the empiricist I have in mind denies this. Kripke showed that there are *contingent* propositions that can be known a priori. One example is *Water is the watery stuff of our acquaintance*. ('Watery' is David Chalmers's nice term for the properties of water by which folk identify it.) All the examples Kripke gave are of propositions that are, to use Gareth Evans's term, deeply necessary (Evans 1979). It is a matter of controversy presently just how to analyse Evans's concepts of deep necessity and contingency, but most of the controversies are over details that are not important right here. I'll simply adopt Stephen Yablo's recent suggestion: a proposition is deeply contingent if it could have *turned out* to be true, and could have *turned out* to be false (Yablo 2002)¹. Kripke did not provide examples of any *deeply* contingent propositions knowable a priori, though nothing he showed rules out their existence.

^{*} This paper has been presented at Cornell University and the Inland Northwest Philosophy Conference, and each time I received valuable feedback. Thanks also to David Chalmers, Harold Hodes, Nicholas Sturgeon, a reader for *Oxford Studies in Epistemology* and, especially, Tamar Szabó Gendler for very helpful comments on various drafts of the paper.

¹ If you prefer the 'two-dimensional' way of talking, a deeply contingent proposition is one that is true in some possible world 'considered as actual'. See (Chalmers forthcoming) for a thorough discussion of ways to interpret this phrase, and the broader notion of so-called 'deep' contingency. Nothing that goes on here will turn on any of the fine distinctions made in that debate - the relevant propositions will be deeply contingent in every plausible sense.

The final debate is a version of **internalism vs externalism** about epistemic justification. The internalist I have in mind endorses a very weak kind of access internalism. Say that a class of properties (intuitively, a determinate) is *introspective* iff any beliefs an agent has about which property in the class (which determinate) she instantiates are guaranteed to not be too badly mistaken.² (Since ‘too badly’ is vague, ‘introspective’ will be vague too, but as we’ll see this won’t matter to the main argument.) My internalist believes the following two claims:

- Which propositions an agent can justifiably believe supervenes in which introspective properties she instantiates, and this is knowable a priori.³
- There exist some introspective properties and some deeply contingent propositions about the future such that it’s a priori that whoever instantiates those properties can justifiably believe those propositions.

My externalist denies one or other of these claims. Typically, she holds that no matter what introspective properties you have, unless some external condition is satisfied (such as the reliability of the connection between instantiating those properties and the world being the way you believe it is) you lack justification. Alternatively, she holds that the connection between introspective properties and justification is always a posteriori. (Or, of course, she might deny both.)

My argument will be that the combination of anti-scepticism, empiricism and internalism is untenable. Since there’s quite a bit to be said for each of these claims individually, that their combination is untenable means we are stuck with a fairly hard choice: accept scepticism, or rationalism, or externalism. Of the three, it *may* seem that externalism is the best, but given how weak the version of

² That a property is introspective does not mean that whenever a subject instantiates it she is in a position to form a not too badly mistaken belief about it. Even if the subject instantiates the property she may not possess sufficient concepts in order to have beliefs about it. And even if she has the concept she may simply have more pressing cognitive needs than forming certain kinds of belief. Many agents have no beliefs about the smell in their ordinary environment much of the time, for example, and this does not show that phenomenal smell properties are not introspective. All that is required is that if she has any beliefs at all about which determinate she instantiates, the beliefs are immune to massive error.

³ There is a delicate ambiguity in this expression to which a referee drew my attention. The intended meaning is that for any two agents who instantiate the same introspective properties, belief in the same propositions is justified. What’s not intended is that if there’s an agent who justifiably believes p , and the introspective properties they instantiate are F_1, \dots, F_n , then any agent who instantiates F_1, \dots, F_n is justified in believing p . For there might be some other introspective property F_{n+1} they instantiate that justifies belief in q , and q might be a defeater for p . The ‘unintended’ claim would be a very strong, and very implausible, claim about the subvenient basis for justification.

internalism is that I'm using, I think we should take the rationalist option seriously.⁴ In this paper I'll just argue against the combination of anti-scepticism, empiricism and internalism, and leave it to the reader to judge which of the three to reject.

Very roughly, the argument for the trilemma will be as follows. There are some propositions q such that these three claims are true.

- (2) If anti-scepticism is true, then I either know q a priori or a posteriori.
- (3) If internalism and empiricism is true, I do not know q a priori.⁵
- (4) If internalism is true, I do not know q a posteriori.

Much of the paper will be spent giving us the resources to find, and state, such a q , but to a first approximation, think of q as being a proposition like *I am not a brain-in-a-vat whose experiences are as if they were a normal person*.⁶ The important features of q are that (a) it is entailed by propositions we take ourselves to know, (b) it is possibly false and (c) if something is evidence for it, then any evidence is evidence for it. I will claim that by looking at propositions like this, propositions that say in effect that I am not being misled in a certain way, it is possible to find a value for q such that (2), (3) and (4) are all true. From that it follows that

For most of the paper I will assume that internalism and anti-scepticism are true, and use those hypotheses to derive rationalism. The paper will conclude with a detailed look at the role internalism plays in the argument, and this will give us some sense of what an anti-sceptical empiricist externalism may look like.

1. A Sceptical Argument

As mentioned, among the many things I know about future, one of the firmest is (1).

- (1) It will snow in Ithaca next winter.

I know this on the basis of inductive evidence about the length of meteorological cycles and the recent history of Ithaca in winter. The inductive sceptic now raises the spectre of Winter Wonderland, a kind of

⁴ Rationalism is supported by BonJour (1997) and Hawthorne (2002), and my argument owes a lot to each of their discussions.

⁵ Aesthetically it would be preferable to have the antecedent of this claim be just that empiricism is true, but unfortunately this does not seem to be possible.

⁶ I.e. I am not a brain-in-a-vat* in the sense of Cohen (1999)

world that usually has the same meteorological cycles as ours, and has the same history, but in which it is sunny every day in Ithaca next winter.⁷ She says that to know (1) we must know that (5) is false, and we do not.

(5) I am living in Winter Wonderland.

Just how does reflection (5) affect my confidence that I know (1)? The sceptic might just appeal to the intuition that I don't know that (5) is false. But I don't think I have that intuition, and if I do it is much weaker than my intuition that I know (1) and that I can infer (5) from (1). James Pryor (2000: 527-8) has suggested the sceptic is better off using (5) in the following interesting argument.⁸

Sceptical Argument 1

- (6) Either you don't know you're not living in Winter Wonderland; or, if you do know that, it's because that knowledge rests in part on your inductive knowledge that it will snow in Ithaca next winter.
- (7) If you're to know (1) on the basis of certain experiences or grounds *e*, then for every *q* which is "bad" relative to *e* and (1), you have to be in a position to know *q* to be false in a non-question-begging way—i.e., you have to be in a position to know *q* to be false antecedently to knowing that it will snow next winter on the basis of *e*.
- (8) (5) is "bad" relative to any course of experience *e* and (1).
- C. You can't know (1), that it will snow next winter on the basis of your current experiences.

An alternative hypothesis *q* is "bad" in the sense used here iff (to quote Pryor) "it has the special features that characterise the sceptic's scenarios—whatever those features turn out to be." (527) To a first approximation, *q* is bad relative to *p* and *e* iff you're meant to be able to know *p* on the basis of *e*, but *q* is apparently compatible with *e*, even though it is not compatible with *p*.

Pryor argues that the best response to the external world sceptic is **dogmatism**. On this theory you can know *p* on the basis of *e* even though you have no prior reason to rule out alternatives to *p*

⁷ If she is convinced that there is no possible world with the *same* history as ours and no snow in Ithaca next winter, the sceptic will change her story so Winter Wonderland's past differs imperceptibly from the past in our world. She doesn't think this issue is particularly relevant to the *epistemological* debate, no matter how interesting the scientific and metaphysical issues may be, and I agree with her.

⁸ Pryor is discussing the external world sceptic, not the inductive sceptic, so the premises here are a little different to those he provides.

compatible with *e*. Pryor only defends the dogmatic response to the external world sceptic, but it's worth considering the dogmatist response to inductive scepticism. According to this response, I *can* come to know I'm not in Winter Wonderland on the basis of my experiences to date, even though I didn't know this a priori. So dogmatism is a version of empiricism, and it endorses (6).⁹ The false premise in this argument, according to the dogmatist, is (7). We can know it will snow even though the Winter Wonderland hypothesis is bad relative to this conclusion and our actual evidence, and we have no prior way to exclude it.

Pryor notes that the sceptic could offer a similar argument concerning justification, and the dogmatist offers a similar response.

Sceptical Argument 2

- (9) Either you're not justified in believing that you're not in Winter Wonderland; or, if you are justified in believing this, it's because that justification rests in part on your justified belief that it will snow in Ithaca next winter.
- (10) If you're to have justification for believing (1) on the basis of certain experiences or grounds *e*, then for every *q* which is "bad" relative to *e* and (1), you have to have antecedent justification for believing *q* to be false—justification which doesn't rest on or presuppose any *e*-based justification you may have for believing (1).
- (11) (5) is "bad" relative to any course of experience *e* you could have and (1).
- C. You can't justifiably believe it will snow in Ithaca next winter on the basis of past experiences.

The dogmatist rejects (10), just as she rejects (7). I shall spend most of my time in the next two sections arguing for (10), returning to (7) only at the end. For it seems there are compelling reasons to accept (10), and hold that the problem with this argument is either with (9) or (11).¹⁰

⁹ It is a version of the kind of internalism discussed in footnote 2, since according to the dogmatist seeming to see that *p* can be sufficient justification for belief in *p*. Pryor's preferred version of dogmatism is also internalist in the slightly stronger sense described in the text, but it seems possible that one could be a dogmatist without accepting that internalist thesis. One could accept, for instance, that seeming to see that *p* justifies a belief that *p*, but also think that seeming to see that *q* justifies a belief that *p* iff there is a known reliable connection between *q* and *p*. As I said, even the weaker version of internalism is sufficient to generate a conflict with anti-scepticism and empiricism, provided we just focus on the propositions that can be justifiably believed on the basis of introspective properties.

¹⁰ Just which is wrong then? That depends on how "bad" is defined. On our final definition (8) will fail, but there are other sceptical arguments, using other sceptical hypotheses, on which (6) fails.

2. Dominance Arguments

The primary argument for (10) will turn on a dominance principle: if you will be in a position to justifiably believe p whatever evidence you get, and you know this, then you are now justified in believing p . This kind of reasoning is perfectly familiar in decision theory: if you know that one of n states obtains, and you know that in each of those states you should do X rather than Y, then you know now (or at least you should know) that you should do X rather than Y. This is a very plausible principle, and equivalent epistemic principles are just as viable. Dominance reasoning can directly support (10) and hence indirectly support (7). (As Vann McGee 1999 showed, the dominance principle in decision theory has to be qualified for certain kinds of agents with unbounded utility functions who are faced with a decision tree with infinitely many branches. Such qualifications do not seem at all relevant here.)

It will be useful to start with an unsound argument for (10), because although this argument is unsound, it fails in an instructive way. Before I can present the argument I need to make an attempt at formalising Pryor's concept of badness.

q is **bad** relative to e and $p =_{df} q$ is deeply contingent, you know p entails $\sim q$, and for any possible evidence e' (that you could have had at the time your total evidence is actually e) there exists a p' such that you know p' entails $\sim q$ and you are justified in believing p' on the basis of e' if e' is your total evidence.

Roughly, the idea is that a bad proposition is one that would be justifiably ruled out by any evidence, despite the fact that it could turn out to be true.¹¹ Using this definition we can present an argument for rationalism. The argument will use some fairly general premises connecting justification, evidence and badness. If we were just interested in this case we could replace q with (5), the proposition that I'm living in Winter Wonderland, r with the proposition that (5) is false, e with my current evidence, and e' with some evidence that would undermine my belief that (5) is false, if such evidence could exist. The intuitions behind the argument may be clearer if you make those substitutions when reading through the argument. But because the premises are interesting beyond their application to this case, I will present the argument in its more general form.

¹¹ Note that there's a subtle shift here in our conception of badness. Previously we said that bad propositions are those you allegedly know on the basis of your actual evidence (if you know p) even though they are logically consistent with that evidence. Now we say that they are propositions you could rule out on *any* evidence, even though they are consistent with your actual total evidence. This is a somewhat narrower class of proposition, but focussing on it strengthens the sceptic's case appreciably.

Rationalist Argument 1

- (12) If you are justified in believing (1), that it will snow in Ithaca next winter, on the basis of e , and you know (1) entails $\sim(5)$, then you are justified in believing $\sim(5)$ when your evidence is e .
- (13) If you are justified in believing r (at time t) on the basis of e , then there is some other possible evidence e' (that you could have at t) such that you would not be justified in believing r were your total evidence e' .
- (14) If you are justified in believing r , and there is no evidence e such that e is part of your evidence and you are justified in believing r on the basis of e , then you are justified in believing r a priori.¹²
- (15) By definition, q is **bad** relative to e and p iff q is deeply contingent, you know p entails $\sim q$, and for any possible evidence e' (that you could have when your evidence is e) there exists a p' such that you know p' entails $\sim q$ and you are justified in believing p' on the basis of e' if e' is your total evidence.
- (16) So, if q is bad relative to e and (1), and you are justified in believing (1) on the basis of e , then you are justified in believing $\sim q$ a priori.

(The references to times in (13) and (15) is just to emphasise that we are talking about your current evidence, and ways it could be. That you could observe Winter Wonderland next winter doesn't count as a relevant alternative kind of evidence *now*.)

Our conclusion (16) entails (10), since (10) merely required that for every bad proposition relative to e and (1), you have 'antecedent' justification for believing that proposition to be false, while (16) says this justification is a priori. ('Antecedent' justification need not be a priori as long as it arrives before the particular evidence you have for (1). This is why (16) is strictly stronger than (10).) So if (10)

¹² David Chalmers noted that (10) and (11) entail that *I exist* is a priori. He thought this was a bad result, and a sufficient reason to modify these premises. I'm perfectly happy with saying, following Kaplan, that *I exist* is a priori. I don't think this proves rationalism, because I think it's also deeply necessary that I exist. (It's not deeply necessary that Brian exists, but that's no objection to what I just claimed, because it's not deeply necessary that I'm Brian.)

This position is controversial though, so I don't want to rest too much weight on it. If you don't think that *I exist* should be a priori, rewrite (11) so that its conclusion is that you would be justified in believing the material conditional $I \text{ exist} \supset r$ a priori. (Note that since I'm presupposing in the dominance argument that all the salient possibilities are ones in which I have some evidence, and hence exist, it's not surprising that *I exist* has a special status within the theory.)

On a separate point, note that I make no assumptions whatsoever here about what relationship must obtain between a justified belief and the evidence on which it is based. Depending on what the right theory of justification is, that relationship might be entailment or constitution or causation or association or reliable connection or something else or some combination of these. I do assume that a posteriori beliefs are somehow connected to evidence, and if the beliefs are justified this relation is properly called *basing*.

is false then one of these premises must be false. I take (15) to define “bad”, so it cannot be false. Note that given this definition we cannot be certain that (5) is bad. We will return to this point a few times.

Which premise should the dogmatist reject? (12) states a fairly mundane closure principle for justified belief. And (13) follows almost automatically from the notion of ‘basing’. A belief can hardly be based in some particular evidence if any other evidence would support it just as well. This does not mean that such a belief cannot be rationally *caused* by the particular evidence that you have, just that the evidence cannot be the rational *basis* for that belief. The dogmatist objects to (14). There is a *prima facie* argument for (14), but as soon as we set it out we see why the dogmatist is correct to stop us here.

Consider the following argument for (14), which does little more than lay out the intuition (14) is trying to express. Assume r is such that for any possible evidence e , one would be justified in believing r with that evidence. Here’s a way to reason a priori to r . Whatever evidence I get, I will be justified in believing that q . So I’m now justified in believing that r , before I get the evidence. Compare a simple decision problem where there is one unknown variable, and it can one of two values, but whichever value it takes it is better for one to choose X rather than Y. That is sufficient to make it true now that one should choose X rather than Y. Put this way, the argument for (14) is just a familiar dominance argument.

Two flaws with this argument for (14) stand out, each of them arising because of disanalogies with the decision theoretic case.

First, when we apply dominance reasoning in decision theory, we look at cases where it would be better to take X rather than Y in every possible case, *and this is known*. This point is usually not stressed, because it’s usually just assumed in decision theory problems that the players know the consequences of their actions given the value of certain unknown variables. It’s not obviously a good idea to assume this without comment in applications of decision theory, and it’s clearly a bad idea to make the same kind of assumption in epistemology. Nothing in the antecedent of (14) specifies that we can know, let alone know a priori, that if our evidence is e then we are justified in believing r . Even if this is true, even if it is necessarily true, it may not be knowable.

Second, in the decision theory case we presupposed it is known that the variable can take only one of two values. Again, there is nothing in the antecedent of (14) to guarantee the parallel. Even if an agent knows of every possible piece of evidence that if she gets that evidence she will be justified in believing r , she may not be in a position to justifiably conclude r now because she may not know that these are all the possible pieces of evidence. In other words, she can only use dominance reasoning to conclude r if she knows *de dicto*, and not merely *de re*, of every possible body of evidence that it justifies r .

So the quick argument for (14) fails. Still, it only failed because (14) left out two qualifications. If we include those qualifications, and adjust the other premises to preserve validity, the argument will work. To make this adjustment, we need a new definition of badness.

q is **bad** relative to e and $p =_{df}$

- a) q is deeply contingent;
- b) p is known to entail $\sim q$; and
- c) it is knowable a priori that for any possible evidence e' there exists a p' such that p' is known to entail $\sim q$, and one is justified in believing p' on the basis of e' .

The aim still is to find an argument for some claim stronger than (10) in sceptical argument 2. If we can do that, and if as the sceptic suggests (5), the Winter Wonderland hypothesis, really is bad, then the only anti-sceptical response to sceptical argument 2 will be rationalism. So the fact that this looks like a sound argument for a slightly stronger conclusion than (10) is a large step in our argument that anti-scepticism plus internalism entails rationalism. (I omit the references to times from here on.)

Rationalist Argument 2

- (12) If you are justified in believing (1), that it will snow in Ithaca next winter, on the basis of e , and you know (1) entails $\sim(5)$, then you are justified in believing $\sim(5)$ when your evidence is e .
- (10') If you are justified in believing r on the basis of e , then there is some other possible evidence e' such that you would not be justified in believing r were your total evidence e' .
- (17) If you know you are justified in believing r , and you know a priori that there is no evidence e you have such that you are justified in believing r on the basis of e , then you are justified in believing r a priori.¹³
- (18) By definition, q is **bad** relative to e and p iff q is deeply contingent, p is known to entail $\sim q$, and it is knowable a priori that for any possible evidence e' there exists a p' such that p' is known to entail $\sim q$, and one is justified in believing p' on the basis of e' .
- (19) So, if q is bad relative to e and (1), and you are justified in believing (1) on the basis of e , then you are justified in believing $\sim q$ a priori.

This is a sound argument for (19), and hence for (10), but as noted on this definition of “bad” (11) may be false. If the Winter Wonderland hypothesis is to be bad it must be a priori knowable that on any evidence

¹³ Again, if you don't think *I exist* should be a priori, the conclusion should be that *I exist* $\supset r$ is a priori.

whatsoever, you'd be justified in believing it to be false. But as we will now see, although no evidence could justify you in believing the Winter Wonderland hypothesis to be true, it is not at all obvious that you are always justified in believing it is false.

3. Hunting the Bad Proposition

A proposition is bad if it is deeply contingent but if you could justifiably believe it to be false on the basis of your current evidence, you could justifiably believe it to be false a priori. If a bad proposition exists, then we are forced to choose between rationalism and scepticism. To the extent that rationalism is unattractive, scepticism starts to look attractive. I think Pryor is right that this kind of argument tacitly underlies many sceptical arguments. The importance of propositions like (5), the Winter Wonderland hypothesis, is not that it's too hard to know them to be false. The arguments of those who deny closure principles for knowledge notwithstanding, it's very intuitive that it's *easier* to know (5) is false than to know (1), that it will snow in Ithaca next winter, is true. So why does reflection on (5) provide more comfort to the inductive sceptic than reflection on (1)? The contextualist has one answer, that thinking about (5) moves the context to one where sceptical doubts are salient. Pryor's work suggests a more subtle answer. Reflecting on (5) causes us to think about *how* we could come to know it is false, and prima facie it might seem we could not know that a priori or a posteriori. It's that dilemma, and not the mere salience of the Winter Wonderland possibility, that drives the best sceptical argument. But this argument assumes that (5) could not be known to be false on the basis of empirical evidence, i.e. that it is bad. If it is not bad, and nor is any similar proposition, then we can easily deflect the sceptical argument. However, if we assume internalism, we can *construct* a bad proposition.

The prima facie case that (5) is bad (relative to (1) and our current evidence *e* - I omit these relativisations from now on) looks strong. The negation of (5) is (20), where *H* is a proposition that summarises the relevant parts of the history of the world.¹⁴

(20) Either $\sim H$ or it will snow in Ithaca next winter.

Now one may argue that (5) is bad as follows. Either our evidence justifies believing $\sim H$ or it doesn't. If it does, then it clearly justifies believing (20), for $\sim H$ trivially entails it. If it does not, then we are justified in believing *H*, and whenever we are justified believing the world's history is *H*, we can inductively infer that it will snow in Ithaca next winter. The problem with this argument, however, is fairly clear: the step

¹⁴ I assume *H* includes a 'that's all that's relevant clause' to rule out defeaters. That is, it summaries the relevant history of the world *as such*.

from the assumption that we are not justified in believing $\sim H$ to the conclusion we are justified in believing H is a modal fallacy. We might be justified in believing neither H nor its negation. In such a situation, it's not obvious we could justifiably infer (20). So (5) may not be bad.

A suggestion John Hawthorne (2002) makes seems to point to a proposition that is more plausibly bad. Hawthorne argues that disjunctions like (21) are knowable a priori, and this suggests that (22), its negation, is bad.

(21) Either my evidence is not e or it will snow in Ithaca next winter.

(22) My evidence is e and it will not snow in Ithaca next winter.

Hawthorne does not provide a dominance argument that (21) is knowable a priori. Instead he makes a direct appeal to the idea that whatever kinds of inference we can draw now the basis of our evidence we could have drawn prior to getting e as conditional conclusions, conditional on getting e . So if I can now know it will snow in Ithaca next winter, prior to getting e I could have known the material conditional *If my evidence is e , it will snow in Ithaca*, which is equivalent to (21). It's not clear this analogy works, since when we do such hypothetical reasoning we take someone to *know* that our evidence is e , and this may cause some complications. Could we find a dominance argument to use instead? One might be tempted by the following argument.

Rationalist Argument 3

(23) I know a priori that if my evidence is e , then I am justified in believing the second disjunct of (21).

(24) I know a priori that if my evidence is not e , then I am justified in believing the first disjunct of (21)

(25) I know a priori that if I am justified in believing a disjunct of (21) I am justified in believing the disjunction (21).

(26) I know a priori that my evidence is either e or not e .

(26) So, I'm justified a priori in believing (21).

The problem here is the second premise, (24). It's true that if my evidence is not e then the first disjunct of (21) is true. But there's no reason to suppose I am justified in believing any true proposition about my evidence. Timothy Williamson (2000: ch. 8) has argued that the problem with many sceptical arguments is that they assume agents know what their evidence is. I doubt that's really the flaw in sceptical arguments, but it certainly is the flaw in the argument that (22) is bad.

The problem with using (22) is that the argument for its badness relied on quite a strong privileged access thesis: whenever my evidence is not e I am justified in believing it is not. If we can find a weaker privileged access thesis that is true, we will be able to find a proposition similar to (22) that is bad. And the very argument Williamson gives against the thesis that we always know what our evidence is will show us how to find such a thesis.

Williamson proposes a margin-of-error model for certain kinds of knowledge. On this model, X knows that p iff (roughly) p is true in all situations within X 's margin-of-error.¹⁵ The intuitive idea is that all of the possibilities are arranged in some metric space, with the distance between any two worlds being the measure of their similarity with respect to X . Then X knows all the things that are true in all worlds within some sphere centred on the actual world, where the radius of that sphere is given by how accurate she is at forming beliefs.

One might think this would lead to the principle B: $p \rightarrow K \sim K \sim p$, that is, if p is true then X knows that she does not know $\sim p$. Or, slightly more colloquially, if p is true then X knows that for all she knows p is true. (I use K here as a modal operator. KA means that X , the salient subject, knows that A .) On a margin-of-error model $p \rightarrow K \sim K \sim p$ is false only if p is actually true and there is a nearby (i.e. within the margin-of-error) situation where the agent knows $\sim p$. But if *nearby* is symmetric this is impossible, because the truth of p in this situation will rule out the knowability of $\sim p$ in that situation.

As Williamson points out, that quick argument is fallacious, since it relies on a too simplistic margin-of-error model. He proposes a more complicated account: p is known at s iff there is a distance d greater than the margin-of-error and for any situation s' such that the distance between s and s' is less than d , p is true at s' . Given this model, we cannot infer $p \rightarrow K \sim K \sim p$. Indeed, the only distinctive modal principle we can conclude is $Kp \rightarrow p$. However, as Delia Graff (2002) has shown, if we make certain density assumptions on the space of available situations, we can recover the principle (27) within this account.¹⁶

$$(27) \quad p \rightarrow K \sim K K \sim p$$

To express the density assumption, let $d(s_1, s_2)$ be the 'distance' between s_1 and s_2 , and m the margin-of-error. The assumption then is that there is a $k > 1$ such that for any s_1, s_2 such that

¹⁵ There's a considerable amount of idealisation here. What's really true is that X is in a position to know anything true in all situations within her margin-of-error. Since we're working out what is a priori knowable, I'll assume agents are idealised so they know what they are in a position to know. This avoids needless complications we get from multiplying the modalities that are in play.

¹⁶ If we translate K as \Box and $\sim K \sim$ as \Diamond , (24) can be expressed as the modal formula $p \rightarrow \Box \Diamond \Diamond p$.

$d(s_1, s_2) < km$, there is an s_3 such that $d(s_1, s_3) < m$ and $d(s_3, s_2) < m$. And this will be made true if there is some epistemic situation roughly ‘half-way’ between s_1 and s_2 .¹⁷ That is, all we have to assume to recover (27) within the margin-of-error model is that the space of possible epistemic situations is suitably dense. Since the margin-of-error model, and Graff’s density assumption, are both appropriate for introspective knowledge, (27) is true when p is a proposition about the agent’s own knowledge.

To build the bad proposition now, let G be a quite general property of evidence, one that is satisfied by everyone with a reasonable acquaintance with Ithaca’s weather patterns, but still precise enough that it is a priori that everyone whose evidence is G is justified in believing it will snow in Ithaca next winter. The internalist, remember, is committed to such a G existing and it being an introspective property. Now consider the following proposition, which I shall argue is bad.¹⁸

(28) I know that I know my evidence is G , and it will not snow in Ithaca next winter.

The negation of (28) is (29).

(29) It will snow in Ithaca next winter, or I don’t know that I know my evidence is G .

It might be more intuitive to read (29) as the material conditional (29a), though since English conditionals aren’t material conditionals this seems potentially misleading.

(29a) If I know that I know that my evidence is G , then it will snow in Ithaca next winter.

To avoid confusions due to the behaviour of conditionals, I’ll focus on the disjunction (29). Assume for now that the margin-of-error model is appropriate for propositions about my own evidence. I will return below to the plausibility of this assumption. This assumption implies that principle (27) is always correct when p is a proposition about my evidence. Given this, we can prove (28) is bad. Note that all my possible evidential states either are, or are not, G . If they are G then by hypothesis I am justified in believing that it will snow in Ithaca next winter and hence I am justified in believing (29). If they are not,

¹⁷ Graff actually gives a slightly stronger principle than this, but this principle is sufficient for her purposes, and since it is weaker than Graff’s, it is a little more plausible. But the underlying idea here, that we can get strong modal principles out of margin-of-error models by making plausible assumptions about density, is taken without amendment from her paper.

¹⁸ If you preferred the amended version of (11) discussed in footnote 12, the bad proposition is *I don’t exist or (28) is true*.

then by the principle (27) I know that I don't know that I know my evidence is G , so I can come to know (29), so I am justified in believing (29). So either way I am justified in believing (29). It's worth noting that at no point here did I assume that I knew whether my evidence was G , though I do assume that I know that having evidence that is G justifies belief in snow next winter.

All of this assumes the margin-of-error model looks appropriate for introspective properties. If it isn't, then we can't assume that (27) is true when p is a proposition about the introspective properties I satisfy, and hence the argument that (29) is knowable a priori fails. There's one striking problem with assuming a priori that we can use the margin-of-error model in all situations. It is assumed (roughly) that anything that is true in all possibilities within a certain sphere with the subject's beliefs at the centre is known. This sphere must include the actual situation, or some propositions that are actually false may be true throughout the sphere. Since for propositions concerning non-introspective properties there is no limit to how badly wrong the subject can be, we cannot set any limits a priori to the size of the sphere. So a priori the only margin-of-error model we can safely use is the sceptical model that says the subject knows that p iff p is true in all situations. For introspective properties the margin-of-error can be limited, because it is constitutive of introspective properties that the speakers beliefs about whether they possess these properties are not too far from actuality. So there seems to be no problem with using Williamson's nice model as long as we restrict our attention to introspective properties.¹⁹

If belief in (29) can be justified a priori, and it is true, does that mean it is knowable a priori? If we want to respect Gettier intuitions, then we must not argue directly that since our belief in (29) is justified, and it is true, then we know it. Still, being justified and true is not irrelevant to being known. I assume here, far from originally, that it is a reasonable *presumption* that any justified true belief is an item of knowledge. This presumption can be defeated, if the belief is inferred from a false premise, or if the justification would vanish should the subject acquire some evidence she should have acquired, or if there is a very similar situation in which the belief is false, but it is a reasonable presumption. Unless we really are in some sceptical scenario, there is no "defeater" that prevents our belief in (29) being an item of knowledge. We certainly did not infer it from a false premise, there is no evidence we *could* get that would undermine it, and situations in which it is false are very far from actuality.

¹⁹ There's a possible complication here related to the point made in footnote 3 about the different ways of formulating the internalist claim. Even if internalism is true, it might be possible for an agent to be radically mistaken about the state of her evidence. For she might think internalism is false, that some extrospective property F is evidentially relevant, and be as mistaken as can be as whether she instantiates F . By assuming that internalism is a priori knowable, we avoid that problem. For the agents we are discussing here are, as mentioned in footnote 15, convenient idealisations who are aware of a priori facts like the truth of internalism.

Since there are no such defeaters, it is reasonable to infer we can *know* (29) a priori. The important premises grounding this inference are an anti-sceptical premise, that we can know (1) on the basis of our current evidence, and the internalist premise that we used several times in the above argument. This completes the argument that the combination of empiricism, internalism and anti-scepticism is untenable.

4. How Externalism Helps

It should be obvious how the rationalist can respond to the above argument - by simply accepting the conclusion. Ultimately I think that's the best response to this argument. As Hawthorne notes, rationalism is the natural position for fallibilists about knowledge to take, for it is just the view that we can know something a priori even though we could turn out to be wrong. In other words, it's just fallibilism about a priori knowledge. Since fallibilism about a posteriori knowledge seems true, and there's little reason to think fallibilism about the a priori would be false if fallibilism about the a posteriori is true, the rationalist's position is much stronger than many have assumed.²⁰ The inductive sceptic also has an easy response - reject the initial premise that in my current situation I know that it will snow in Ithaca next winter. There are other responses that deserve closer attention: first, the inductive sceptic who is not a universal sceptic, and in particular is not a sceptic about perception, and second the externalist.

I said at the start that the argument generalises to most kinds of scepticism. One kind of theorist, the inductive sceptic who thinks we can nonetheless acquire knowledge through perception, may think that the argument does not touch the kind of anti-sceptical, internalist, empiricist position she adopts. The kind of theorist I have in mind says that the objects and facts we perceive are constitutive of the evidence we receive. So given we are getting the evidence we are actually getting, these objects must exist and those facts must be true. She says that if I'd started with (30), instead of (1), my argument would have ended up claiming that (31) is bad for some *G*.

(30) A hand exists.

(31) A hand exists, or I don't know that I know that I'm perceiving a hand.

She then says that (31) is not deeply contingent, since in any situation where the first disjunct is false the second is true, so it cannot be bad. This response is correct as far as it goes, but it does not go far enough to deserve the name anti-sceptical. For it did not matter to the above argument, or to this response that (1)

²⁰ As Bonjour points out, rationalism has fallen into such disrepute that many authors leave it out even of surveys of the options. This seems unwarranted given the close connection between rationalism and the very plausible thesis of fallibilism.

is about the future. All that mattered was that (1) was not *entailed* by our evidence. So had (1) been a proposition about the present that we cannot directly perceive, such as that it is not snowing in Sydney *right now*, the rest of the argument would have been unaffected. The summary here is that if one is suitably externalist about perception, so one thinks the existence of perceptual states entail the existence of the things being perceived, one can accept this argument, accept internalism, accept empiricism, and not be an *external world* sceptic. For it is consistent with such a position that one know the existence of the things one perceives. But on this picture one can know very little beyond that, so for most practical purposes, the position is still a sceptical one.

The externalist response is more interesting. Or, to be more precise, the externalist responses are more interesting. Although I have appealed to internalism a couple of times in the above argument, it might not be so clear how the externalist can respond. Indeed, it may be worried that by exercising a little more care in various places I could have shown that everyone must accept either rationalism or scepticism. That is the conclusion Hawthorne derives in his paper on deeply contingent a priori knowledge, though as noted above he uses somewhat more contentious reasoning than I do in order to get there. To conclude, I will argue that the internalism is crucial to the argument I have presented, and I will spell out how the externalist can get out of the trap I've set above.

One easy move that's available to an externalist is to deny that any facts about justification are a priori. That blocks the move that says we can find a *G* such that it's a priori that anyone whose evidence is *G* can know that it will snow in Ithaca next year. This is not an essential feature of externalism. One can be an externalist about justification and still think it is a priori that if one's evidence has the property *is reliably correlated with snow in the near future* then it justifies belief that it will shortly snow. But the position that all facts about justification are a posteriori fits well with a certain kind of naturalist attitude, and people with that attitude will find it easy to block the sceptical argument I've presented.

Can, however, we use an argument like mine to argue against an anti-sceptic empiricist externalist who thinks some of the facts about justification *can* be discovered a priori? The strategy I've used to build the argument is fairly transparent: find a disjunctive a priori knowable proposition by partitioning the possible evidence states into a small class, and adding a disjunct for every cell of the partition. In every case, the disjunct that is added is one that is known to be known given that evidence. If one of the items of knowledge is ampliative, if it goes beyond the evidence, then it is possible the disjunction will be deeply contingent. But the disjunction is known no matter what.

If internalism is true, then the partition can divide up evidential states according to the introspective properties of the subject. If externalism is true, then such a partition may not be that *useful*, because we cannot infer much about what the subject is justified in believing from the introspective properties she instantiates. Consider, for example, the above partition of subjects into the *G* and the not-*G*,

where G is some introspective property, intuitively one somewhat connected with it snowing in Ithaca next year. The subjects that are not- G know that they don't know they know they are G , because they aren't. Externalists need not object to this stage of the argument. They can, and should, accept that a margin-of-error model is appropriate for introspective properties. Since it's part of the nature of introspective properties that we can't be *too* badly wrong about which ones we instantiate, we're guaranteed to satisfy some reliability clause, so there's no ground there to deny the privileged access principle I defended above.

The problem is what to say about the cases where the subject is G . Externalists should say that some such subjects are justified in believing it will snow in Ithaca next winter, and some are not. For simplicity, I'll call the first group the reliable ones and the others the unreliable ones. If I'm G and reliable, then I'm justified in believing it will snow, and hence in believing (29). But if I'm G and unreliable, then I'm not justified in believing this. Indeed, if I'm G and unreliable, there is no obvious argument that I'm justified in believing *either* of the disjuncts of (29). Since this is a possible evidential state, externalists should think there is no dominance argument that (29) is a priori knowable.

Could we solve this by adding another disjunct, one that is guaranteed to be known if I'm G and unreliable? There is no reason to believe we could. If we're unreliable, there is no guarantee that we will *know* we are unreliable. Indeed, we may well believe we are reliable. So there's no proposition we can add to our long disjunction while saying to ourselves, "In the case where the subject is G and unreliable, she can justifiably believe *this* disjunct." If the subject is unreliable, she may not have *any* justified beliefs about the external world. But this is just to say the above recipe for constructing bad propositions breaks down. Externalists should have no fear that anything like this approach could be used to construct a proposition they should find bad. This is obviously not a positive argument that anti-sceptical empiricist externalism is tenable, but it does suggest that such a position is immune to the kind of argument I have presented here.

References

- Chalmers, David (forthcoming) "The Foundations of Two-Dimensional Semantics" *Philosophical Studies*, forthcoming.
- Cohen, Stewart (1999) "Contextualism, Skepticism, and the Structure of Reasons" *Philosophical Perspectives* 13: 57-89.
- Evans, Gareth (1979) "Reference and Contingency" *The Monist* 62: 161-89.
- Graff, Delia (2002) "An Anti-Epistemicist Consequence of Margin for Error Semantics for Knowledge" *Philosophy and Phenomenological Research* 64: 127-142.
- Hawthorne, John (2002) "Deeply Contingent A Priori Knowledge," *Philosophy and Phenomenological Research*, 64: 247-269.
- Jackson, Frank (1998) *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.
- McGee, Vann (1999) "An Airtight Dutch Book" *Analysis* 59:257-265.
- Pryor, James (2000) "The Skeptic and the Dogmatist" *Noûs* 34: 517-49
- Williamson, Timothy (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.
- Yablo, Stephen (2002) "Coulda, Woulda, Shoulda" in Tamar Szabó Gendler and John Hawthorne (eds) *Conceivability and Possibility*. Oxford: Oxford University Press, pp. 441-492.