

APPENDIX A: PRYOR AND BYRNE'S COMPARISONS

Some who complain that AI is a weak argument due to the weakness of its first premise have other arguments that they are seeking to comparatively promote as more powerful than AI. James Pryor offers the skeptic what he takes to be a better argument, and Alex Byrne comparatively praises other arguments for very counter-intuitive conclusions (that are not examples of epistemological skepticism). It may be instructive for me to explain why I reject these comparative claims.

PRYOR

Pryor registers this complaint against AI's first premise:

Some philosophers refuse to allow the skeptic to use claims like 'I can't know I'm not being deceived' as *premises* in his reasoning. Maybe skeptical argument can *convince us* that we can't know we're not being deceived; but why should we grant such a claim as *a premise* in a skeptical argument?' (Pryor 2000: 522).

So Pryor offers the skeptic what he thinks is a better argument. Pryor's proposed argument can be viewed as a descendant of AI—and a descendant that is in ways more subtle than is the original, trying to make explicit more of what one might suppose is the implicit reasoning by which thoughts of skeptical hypotheses can inspire skeptical conclusions. But increased subtlety has its dangers when trying to produce a powerful argument. Especially when the premises of the original, less subtle argument are as intuitively powerful as AI's are, a more subtle replacement argument can easily fail to capture the intuitive power of the original. And that is what I think happens here.

Pryor replaces the AI's first premise with:

(5) Either you don't know you're not being deceived by an evil demon; or, if you do know you're not being deceived, it's because that knowledge rests in part on things you know by perception. (Pryor 2000: 524)

Pryor prefers to work with the evil genius, rather than the BIV hypothesis. I assume that makes no significant difference to our evaluation of these arguments. The more relevant difference here is that Pryor's (5) is in an important way weaker than AI's (1). Accordingly, (5) is at least as plausible as (1). And I'd even agree that it's a bit *more* plausible than (1). But, at least for me, *just* a bit. We can break this matter down into two questions:

a) Do we know that not-H because this knowledge is based in part on things we know by perception?

b) Do we know that not-H in some other way?

In putting (1) forward as a premise, my skeptic is claiming, and asking us to agree that, the answer to both of these questions is 'no.' In putting (5) forward, Pryor's skeptic is only claiming that the answer to (b) is 'no.'

For reasons we will soon see, the power of the skeptical arguments we are comparing does not really fully rest on how (5) compares with (1). Still, for what it's worth, here's how the comparison intuitively strikes me. When I consider whether I know that skeptical hypotheses like the BIV or the evil demon hypotheses are false, one of the first thoughts to pop into my head is a question: 'How *could* I *possibly* know something like that?' Having asked that question, I immediately start fishing around for a possible answer—some potential way that I might have the exotic knowledge in question, and one of my very next thoughts is something along the lines of 'Certainly *not* by basing it on something like that I have hands, or that I'm sitting at a desk, or the like. Those perceptual beliefs can't be used to support a conclusion like that I'm not a BIV or not a victim of an evil demon. So if I am to know such things as that I'm not a BIV at all, it must be in some other way. But what other way?' Then I start fishing around for some other possible way. And some possibilities don't seem all that bad in themselves, and certainly not as bad as the thought that I might successfully base the conclusion on my perceptual beliefs, but by this point, admittedly, my thought has gotten fairly philosophical, and has gone well beyond the intuitive. (I explain my own answer to this question in Chapter 7, where I also compare it at length with a Pryor-like account on which we know that radical skeptical hypotheses by basing such knowledge on perceptual judgments.) But there is always the fairly intuitive thought that perhaps I might know such exotic things in some way I haven't thought of yet. Indeed, 'I just know it—don't ask me how!', while certainly far from a great answer, perhaps because of its sheer

openness strikes me as intuitively more promising than ‘I successfully base it in large part on my perceptual beliefs.’ So it has always intuitively struck me that the prospects for a positive answer to (b) are brighter than for (a). Which is why, to me, (5) doesn’t seem much more plausible than (1). (But again, see Chapter 7 for a more thorough comparison.)

It’s clear that this matter strikes Pryor very differently. He admits (2000: 524) that some philosophers reject (5), but reasonably doesn’t let that disqualify (5) from being a legitimate premise. This is reasonable because, after all, unanimous agreement among philosophers, or even being relatively uncontroversial among philosophers, is way too high a standard—which is something we should keep in mind when assessing (1)’s premise-worthiness, as well. But why is (5) so much better as a premise than (1)? In defense of (5), Pryor writes:

I think that this [(5)] is a very plausible premise. The hypothesis that all of our present experiences are the deceptions of an evil demon is not absurd. It seems to be a genuine metaphysical possibility. So we can’t reject that hypothesis out of hand. If we do know that we’re not being deceived by an evil demon, it’s plausible that that knowledge would have to rest on things we know about our environment on the basis of perception. (2000: 524)

The first sentence above states the judgment that Pryor is here seeking to support—that (5) is very plausible—which I assume means something strong enough to imply that it’s significantly more plausible than (1), toward which he offers no such praise. I am largely in agreement with the next two sentences: Although the evil demon hypothesis certainly can strike many in some ways as ‘absurd’ (and, as reported in Chapter 2, I incorporate such reactions into my account in Chapter 7 of how we come to know by ordinary standards that we’re not BIVs), I suspect that by ‘absurd’, Pryor here means something like ‘metaphysically impossible,’ in which case his second and third sentences amount to the same thing, and I agree with both. But I don’t find this point of agreement pushing me toward seeing why (5) is so plausible in a way that (1) is not. Even while keeping the thought that the skeptic’s hypotheses are metaphysically possible in the forefront of my mind, I’m still just about as strongly inclined to think that if I’m to know at all that I’m not in the skeptic’s scenario, it will have to be in some way other than by basing that judgment on my perceptual beliefs. Since the fifth and final sentence of the passage seems just a repetition of the conclusion of the reasoning, a lot seems to rest on the claim expressed in the fourth sentence. And here I must admit that I’m not all that clear on

what Pryor means by rejecting a claim ‘out of hand.’ Given what the following sentence says, my suspicion is that one can only reject a claim ‘out of hand’ in the intended sense if you know *a priori* that it is false. In that case, what the argument of this passage really turns on is a commitment of Pryor’s to some (time-honored) principle like ‘No *a priori* knowledge of deeply contingent truths.’¹ As we see in Chapter 7, I don’t (at all) share this commitment, and I will defend our having *a priori* justification for and *a priori* knowledge of the fact that various skeptical hypotheses are false. So, I disagree with that key part of Pryor’s case even on the reading of ‘out of hand’ just considered, but in any case, rejecting the skeptic’s hypothesis ‘out of hand’ *sounds* at least somewhat like some of the potential ways of coming to know that I’m not the victim of an evil demon nor a BIV that seem to me a bit more promising than by basing that belief on my perceptual beliefs. If there *are* some things that I can come to know by in some sense rejecting their negations ‘out of hand,’ that I’m not a BIV seems to me like a pretty good candidate. After all, as I’ve said, its negation, that I am a BIV, does strike me as in *some* potentially relevant way (though not the way I’m inclined to read Pryor’s use of the word) as ‘absurd.’ Anyway, I’m not finding much here to convince me that basing the denials of our skeptical hypotheses on my perceptual beliefs is the best hope for coming to know that the hypotheses are false.

Still, Pryor’s (5) is of course, even to my thinking, a bit more plausible than is AI’s (1). And Pryor, who no doubt finds the difference in plausibility greater than I do (and apparently *much* greater than I do), stresses this greater plausibility when urging that the whole skeptical argument he is offering the skeptic is better, as for instance, in this key wrapping-the-matter-up statement where Pryor compares his argument with AI:

I believe that this argument reconstructs the skeptic’s reasoning more accurately than the argument (1)-(3) [AI]. In addition, since this new argument does not rely on a bald claim like

¹ Well, the principle is time-honored without the limitation to the *deeply* contingent. The distinction between the deeply contingent and the superficially contingent is relatively recent and due to (Evans 1979: 161), who, responding to examples of Kripke, acknowledges that there can be *a priori* knowledge of superficially contingent truths, but who still insists that ‘it would be intolerable for there to be a statement which is both knowable *a priori* and deeply contingent’ (1979: 161). Since we will not be dealing with any examples of merely superficially contingent claims, we will not explore this distinction here.

(1), but rather on the more plausible (5), it poses a compelling and formidable threat to our possession of perceptual knowledge. (Pryor 2000: 529)

But here is why the matter of how much plausibility Pryor gains for the skeptic by weakening (1) to (5) is not all-important to our comparison of the arguments: The power of the arguments in question, of course, depends on the plausibility of *all* their premises. So, we should compare weakest link with weakest link: Even if (5) is the premise in Pryor's argument that most resembles (1) (we might call these the 'base premises' of their respective arguments), if his weakening of the skeptic's base premise, from (1) to (5), causes him to have to strengthen other parts of the skeptical argument in ways that make some other part become his new argument's weak link, then what he is offering the skeptic might not be such a great deal in the end (even if, against my own judgment, the difference in plausibility between (5) and (1) is quite significant), and what we should be most interested in comparing to (1) in terms of plausibility is the new argument's weakest link. And of course we should also consider whether one argument has more at least fairly weak links than the other.

Pryor combines (5) with these two premises to form his new skeptical argument:

SPK: If you're to know a proposition p on the basis of certain experiences or ground E , then for every q which is 'bad' relative to E and p , 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 p on the basis of E .

(6) The hypothesis that you're being deceived by an evil demon is 'bad' relative to any course of experience E and perceptual belief p . (Pryor 2000: 528)

To judge the plausibility of these premises, one needs to know, for one thing, what is meant by Pryor's technical use of 'bad.' The explanation begins as follows:

Say that an alternative to p is a 'bad' alternative just in case it has the special features that characterize the skeptic's scenarios—whatever those features turn out to be. Different skeptical arguments will rely on different accounts of what makes an alternative 'bad.' Here are some examples. (2000: 527)

And I do indeed find the examples that Pryor goes on to give helpful in getting some decent grip on what Pryor means by 'bad.' We're then told:

For our discussion it does not matter which of these accounts of ‘badness’ the skeptic adopts. Let’s suppose the skeptic does have some such account. His argument would then go as follows. (2000: 527)

How shall we judge (6), which says simply says that the evil demon hypothesis is ‘bad’, without a more precise understanding of ‘bad’?² Best here to simply resolve to make it a necessary requirement for any account of ‘badness’ to be even minimally adequate that the BIV hypothesis be ‘bad’, and then the question will be whether, in some such sense that makes (6) true, SPK, the ‘Skeptical Principle about Knowledge,’ is true as well. That way, as Pryor writes, ‘Everything will turn on the principle SPK’ (2000: 529).

So understood, then, how plausible *is* SPK? One would think that this is absolutely vital to evaluating the power of the argument Pryor proposes, and for comparing its power with that of AI. And this is especially so for Pryor himself, because, since SPK is the component of the skeptical argument he constructs that he rejects, it looks as if Pryor thinks SPK is the weak link of that argument. But, as we saw, Pryor at least sometimes writes as if the comparison rests solely on how (5)’s plausibility compares with (1)’s. And, tellingly, he treats the plausibility of SPK as very much an open question:

How the skeptic motivates this principle will depend on his choice of skeptical hypothesis and his account of what makes a hypothesis ‘bad.’ Suppose the skeptic does persuade us to accept SPK ... (2000: 529)

And it seems enough that maybe the skeptic might be able to motivate it somehow, sometime:

² I don’t mean to be complaining here on the imprecision of Pryor’s use of ‘bad.’ His use of that term actually seems insightful. And it may seem that I’m in no position to complain, since, in my own treatment of skepticism, though not using the very term ‘bad’, I make use of a notion that plays much the same role when I write of the feature—which I go on to identify as insensitivity—of skeptical hypotheses that makes them effective skeptical weapons. Indeed, in a note Pryor (reasonably) attributes to me an (insensitivity) account of ‘badness’ (Pryor 2000: 544, n. 18). And my account too is imprecise, and needs further sharpening. (For some of the needed sharpening, see Chapter 6.) But for me such a notion comes in as part of my own anti-skeptical deflationary explanation for why we get pulled in by the skeptic’s argument. Pryor’s skeptic has the notion of badness as part of the skeptic’s argument itself. That does complicate the issue of how plausible the premises that contain the technical use of ‘bad’ are in the first place. But I don’t think any significant harm is done, if we, as I suggest in the body of this paragraph, resolve to understand ‘bad’ in a way that makes (6) unproblematically true.

The skeptic still has work to do: he has to persuade us to accept SPK. We have not looked in detail at how he might do that. (2000: 529)

It's interesting to compare this with Pryor's treatment of AI's first premise, which, recall, was: 'Maybe skeptical argument can *convince us* that we can't know we're not being deceived; but why should we grant such a claim as *a premise* in a skeptical argument?' And (1) is thereby rejected as a desirable premise. But there is no such rejection of SPK; the mere thought that the skeptic might be able to motivate it seems enough to get it into the reasonable premise game.

Comparing weakest link with weakest link, at least as Pryor himself would judge that matter, then, we end up weighing (1), which, as I've urged in Chapter 2, is very plausible indeed (and there's some reason to think it *might* be considerably more intuitive still, and perhaps remarkably powerful indeed), against SPK, the plausibility of which Pryor seems reluctant to even vouch for.

Though it looks as if Pryor would classify SPK as the weakest link in the skeptical argument he constructs, I should here briefly address the question of whether *I* so classify it, and register a couple of thoughts that may help readers in their evaluation of it. And that's a very tricky call for me. As I've noted, Pryor's weakening of the content of the skeptic's base premise does strengthen its plausibility, but, for me, not by much. And in the end, I reject (5) in the same contextualist way I reject (1): I deny it where it is evaluated at ordinary standards for knowledge. But what of his skeptic's 'bridge premise', SPK, which I copy here for convenience:

SPK: If you're to know a proposition *p* on the basis of certain experiences or ground E, then for every *q* which is 'bad' relative to E and *p*, 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 *p* on the basis of E (Pryor 2000: 528)?

For me, as I suspect for most, SPK's plausibility turns largely on how we are to understand the 'antecedently' that it contains; so, key here will be Pryor's explanation of knowing one thing 'antecedently' to knowing another / 'epistemic priority' at (Pryor 2000: 524-6). And it's difficult for me to say whether I accept SPK in the way I accept the AI skeptic's bridge premise, (2): whether I accept it as true when evaluated at any given epistemic standard, or even just as evaluated by ordinary standards for knowledge. Given most of what Pryor writes in explanation

of what I agree is that fairly intuitive and important notion, I'm inclined to think, for reasons given in Chapter 7, that neither our simple perceptual beliefs nor our beliefs in the non-obtaining of radical skeptical hypotheses is epistemically downstream from the other, nor do the former *have to be* epistemically downstream from the latter to constitute knowledge. That makes it sound as if I should reject SPK. However, Pryor's discussion of the notion makes me think that what he means here is not a requirement that beliefs that skeptical hypotheses are false be epistemically upstream from simple perceptual beliefs if the latter are to be known, but just that the former be known in a way that does *not* make them epistemically downstream from the latter for the latter to be known, in which case my own inclination is to accept SPK. On this reading, my acceptance of SPK and Pryor's own rejection of it constitutes one of the two, related, big differences in our views. Much more on this in Chapter 7. However, in addition to what I take to be the general direction of Pryor's explanation of epistemic priority, Pryor also mixes in what to my thinking is significantly different talk about whether we have non-question-begging ways of rejecting skeptical hypotheses—and note that such phrasing makes it into Pryor's formulation of SPK itself (see above). And as I'm inclined to understand such phrases, I think we don't have, nor do we need (in order for our perceptual beliefs to constitute knowledge, at least by ordinary standards) non-question-begging ways of rejecting radical skeptical hypotheses. That would suggest I should reject SPK. However, I suspect that the best way to understand Pryor here is to go with what I have called the general direction of his explanation, and then construe him as using 'question-begging' in a way that fits in with that general direction of his thought. All of which ultimately, though quite tentatively, has me thinking that I accept SPK, as I think that's best understood. At any rate, this is for me a very tricky call (on what is probably a fairly important matter), and is likely to be so for others, so readers should definitely take a careful look at (Pryor 2000: 524-6) for themselves. The trickiness may well explain Pryor's seemingly very restrained enthusiasm for the plausibility of SPK—but then makes it difficult, for me at any rate, to understand his relative enthusiasm for the skeptical argument in which it plays an essential role.

Pryor has another reason for preferring the skeptical argument he formulates over AI. His second complaint against AI is that it 'does not generalize in the same ways that the skeptic's reasoning intuitively seems to generalize' (2000: 522). What Pryor claims to intuit here is that there is a 'core structure of skeptical reasoning' which serves to undermine both knowledge and justification. But if you substitute

'am not justified in believing' for 'don't know' throughout AI, the result is not a very powerful argument. This is closely related to Pryor's main worry about AI, which is focused on premise (1), because the problem with trying to aim AI at justification is that that premise is even less plausible when adjusted to being about justification, prompting Pryor to exclaim: 'Why on earth should we accept this premise?' (2000: 523). By contrast, Pryor thinks the argument he offers the skeptic generalizes nicely.

I'm unmoved by this secondary argument of Pryor's for two main reasons. First, while it would be nice if the best skeptical argument aimed at knowledge could be easily converted to a powerful argument against justification as well, it seems that we do the skeptic no favors if we ascribe to her the argument that would be weaker independent of this generalizing consideration on the grounds that generalizing is itself an important factor. The arguments against knowledge should be judged on their own comparative merits, as I've been trying to do here.

But second, I don't think Pryor's argument really is any better at generalizing than is AI, anyway, though I'll be brief about this, both because it isn't that important (given my first reason), but also because the considerations important to comparing the arguments against justification are just analogues of those we've just been through in comparing the arguments against knowledge. First, Pryor is right that (1) is less plausible when it's modified to target justification, but, it seems to me, the same thing happens to his (5). Pryor thinks (5) doesn't lose any plausibility at all in the shift, but the reason he offers is only this: 'This premise just says that our reasons for believing we're not deceived are partly perceptual. Hence, it should be as plausible as its analogue (5) for knowledge was' (2000: 531). But an important alternative to basing the conclusion that one is not being deceived on one's perceptual beliefs becomes more plausible when we're asking just about justification rather than knowledge: That one might be *justified* in believing that H is false via the route of rejecting H 'out of hand' seems more plausible than that one might thereby come to *know* that not-H. Still, as was the case with the originals, the analogue of (5) does seem more plausible than (1), but just a bit, to my thinking.

But the plausibility of SPK seems to take a *big* hit when it is converted to SPJ (the Skeptical Principle about Justification). And as happened with SPK, Pryor's treatment of SPJ doesn't inspire much confidence in the principle. Here is Pryor's evaluation, which seems extremely unenthusiastic (the first sentence in particular is

very far indeed from a ringing endorsement of SPJ's intuitive power!), but also contains a surprising claim comparing SPJ's plausibility with that of SPK:

This principle SPJ is not obviously false. There is some plausibility to the idea that we're entitled to rely on our perceptual beliefs only if we're antecedently entitled to the assumptions that we're not being deceived by an evil demon, or dreaming, and so on. Hence, although the details will depend on the skeptic's account of 'badness,' the major premises in the skeptic's argument (5)-(9) [the skeptical argument against knowledge that Pryor formulates] seem to be just as defensible when translated into a skeptical argument against the possibility of justification. (2000: 531)

Pryor's comparative claim here seems to me mistaken. I've already discussed how (5)'s plausibility seems diminished when it's converted so as to target justification rather than knowledge. The same thing threatens to happen when SPK is converted to SPJ, at least on some ways of understanding the key notions they have in common: That in the relevant non-question-begging, 'antecedent' way one has to be in a position to eliminate every single one of the 'bad' alternatives to the likes of *I have hands*, no matter how outlandish, in order to *know* based on experience that one has hands can seem a pretty bold skeptical claim already, but that one has to be in a position to jump through all of those same outlandish antecedent hoops to even be *justified* (based on experience) in thinking you have hands seems significantly bolder still.³ But Pryor directs no 'Why on earth should we accept that?' exclamation, followed by a questioning of its premise-worthiness, at SPJ. Still, as one might expect given the extremely muted praise for it given above, SPJ in Pryor's hands seems to just roll over and very easily die, with no fretting at all over what seems to do less violence to our epistemic thought, as soon as it's brought into conflict with a real intuition (which ultimately happens at 2000: 537).

To my thinking, then, the skeptical argument against knowledge that Pryor offers the skeptic fails to measure up to AI in terms of its intuitive power. Which is not to say that Pryor's argument isn't interesting in its own right. Being less powerful than AI is not a very damning criticism in my book, since AI seems so

³ Here I should make clear that Pryor's skeptic isn't demanding that one has to believe that all these 'bad' alternatives are false, just that you *have* the right kind of the justification for eliminating them in the right way. That still seems a quite strong skeptical demand, especially when it concerns the mere justification of our perceptual beliefs.

remarkable in terms of generating a strong, surprising conclusion from plausible premises. It is not for nothing that AI got to be the ‘canonical’ skeptical argument.

BYRNE

Alex Byrne argues that AI and other skeptical arguments compare poorly with other paradoxes: While some philosophical paradoxes, like the liar and the sorites, are ‘very hard,’ the skeptical paradoxes are not: ‘The sceptic is just another guy with a bad argument’ (Byrne 2004: 299). But I don’t see how he justifies such a combination of assessments without employing a double-standard. (To keep this manageable, I’ll focus here on the comparison with the sorites, but similar comments would hold for the liar, I believe.)

Byrne’s argument against the hardness of AI is based on the weakness of the skeptic’s first premise (that one doesn’t know that one is not a BIV), which I’ve agreed is the argument’s weak link. At times Byrne seems to be setting the bar this premise must clear to render AI ‘hard’ quite high: He thinks that the ground rules of the right ‘sceptical game’ have it that ‘the sceptic is only allowed to appeal to more-or-less uncontroversial or obvious premises’ (2004: 300). Though my classroom survey results, which we discussed in section 8 of Chapter 2, may suggest a little caution here (at least in terms of intuitive obviousness, though certainly not in terms of being uncontroversial among philosophers), my own inclination is to agree that the skeptic’s first premise doesn’t meet this standard. But surely the same must be said about the weakest link of a sorites argument! At other points, Byrne seems to be setting the bar lower—and still claiming that the skeptic’s first premise doesn’t clear it: At (2004: 303) Byrne writes that the skeptic’s other premise (our (2)) is ‘extremely plausible,’ and indicates that if the skeptic’s first premise is ‘likewise plausible,’ then ‘the canonical sceptical argument presents us with a serious paradox.’ Whether to agree with Byrne about whether (1) clears this lower bar is a close call for me. The skeptic’s first premise is certainly ‘plausible,’ I’d say; but ‘extremely plausible’? Tough call. But in light of our discussion of the power of this premise in sections 7-9 of Chapter 2, it certainly seems that it *can be* even extremely plausible—to many people, in many settings, and certainly not just to skeptically

inclined philosophers. In the end, I'm guessing that whatever difference there is between Byrne and me here likely doesn't run very deep (though again my classroom survey results may suggest we might both do well to exercise some caution here).

But I will then have a rather sharp difference with him on his comparative claim, for I certainly don't think the sorites is any better off at all on this score, much less any kind of clear winner over AI. Indeed, I count it the loser. Now, Byrne doesn't even tell us what form of sorites argument he has in mind. His attention is negatively focused on the skepticism side of the comparison (which I suppose is how the severe double-standard I suspect him as using arose). On one way of running it, a sorites argument will contain a premise *something* like this: For all n , if n grains of sand do not make a heap, then $n+1$ grains do not make a heap. And wouldn't a very similar assessment of such a premise be in order here: that while it is plausible and maybe (though this is pushing it a bit, or maybe more than just a bit) extremely plausible, it is very far from compelling, uncontroversial, or obvious? For my money, I find the skeptic's weak link significantly more plausible than that of the sorites-monger, at least when his argument includes a premise like the one just discussed, but that's up for grabs, I guess.

Above, I construe the sorites as a short argument, the key premise of which is a generalization over a big class of conditionals. It could instead proceed in a long series of steps and inferences, starting with some 'base premise', and then employing a long series of particular conditionals, featuring particular natural numbers: e.g., *If 5 grains of sand do not make a heap, then 6 grains of sand do not make a heap.* (We can understand the conditionals as material conditionals, since they're strong enough for the argument's purposes, and we want these premises to be as weak as possible so far as their content goes, so they can be as strong as possible so far as their plausibility goes.) Anyway, I've always thought that using lots of particular little conditionals just spreads the weakness of the sorites argument out, and results in an argument that has *lots* of premises that are far from being intuitively compelling—the conditionals that concern the cases in the 'gray zone.' Depending on one's views on vagueness, bivalence, etc., one *might* think that each of these premises has a high probability of being true, but they still won't in important ways be intuitively compelling. (Perhaps not that unlike claims to the effect that S1 isn't the winner of the lottery, S2 isn't the winner, etc., where we know exactly one

of the very many of those S's is a winner. Each claim is very likely true, but not so intuitively compelling in an important way.) But one might instead hold, or at least hold open (as possibly right), views on which maybe whole bunches of these little (material) conditionals are neither true nor false, and they don't even have a high probability of being true.⁴ At any rate, no matter one's basic philosophical orientation, I don't think one has to feel that much intellectual discomfort to think of one of these particular conditionals that concerns a very gray area of the relevant spectrum: 'Yeah, that one might be a [or perhaps the] loser here [and if it isn't, one of its neighbors is].'

I'll resist further speculation on what form of sorites 'paradox' might be so frightfully 'hard', except to add the following (which may be helpful in indicating where I'm coming from to some readers). Some ways to try to derive a real problem (an explicit contradiction, or some other intuitive clunker) from sorites reasoning involve an assertion of bivalence. But then, at least to my thinking, you're adding a premise that is very far from compelling (and indeed false)—though many (not me) will find it quite plausible.

Byrne initially supports his enthusiasm for the paradoxes he likes in terms of the sophistication of the attempts to resolve them, and suggests that those that don't measure up (which include the skeptical paradoxes) don't similarly teach us deep things: 'Some paradoxes, for instance the liar and the sorites, are very hard. The defense of a purported solution to either of these two inevitably deploys the latest in high-tech philosophical weaponry. On the other hand, some paradoxes are not at all hard, and may be resolved without much fuss. They do not contain profound lessons about the human condition' (2004: 299). But his argument does not proceed by looking at the 'philosophical weaponry' used in attempts to resolve skeptical puzzles and arguing that it pales in comparison to the weaponry deployed against the sorites; nor does he tell us what 'profound lessons' have been learned via study of the sorites, and how those deep truths about the human condition so learned are so

⁴ If one employs a supervaluational semantics to the (material) conditionals in question, on which they are neither true nor false if they are true on some precisifications of the key vague term they include and false on other precisifications, then one might think the particular conditionals involved go truth-value-less throughout some 'gray zone.'

much deeper and more important than what we can learn from a study of the skeptical paradoxes.

My suspicion is that Byrne's somewhat melodramatic wording here is a reaction against some of the grandiose claims made for the power of skepticism by those who do find the skeptical paradoxes deep. The talk about 'profound truths' about 'the human condition' may well derive in particular from Stroud, who is wont to talk thus, and who is a main target of Byrne's. But then Byrne should also notice that, when in a ho-hum mood, it sounds about as silly to suggest that the sorites has really taught us any profound truths about the human condition!

I am not denying the philosophical importance of the sorites. It need not be any profound paradox for it to be well worth working on. It may be unclear what is the best way of escaping some formulations of the problem, and we may well learn something of philosophical significance (though likely not anything that will deserve to be called a 'profound truth human condition') in figuring out the best way to make the denial of some premise go down as smoothly as possible. But *that* much seems likely in the cards for a study of AI, too—or if it isn't, Byrne at least hasn't told us why we shouldn't be hopeful here.