# THE DOCTOROIDS AND THE PROBLEM OF ACCEPTABLE STARTING POINTS

## 0. Introduction

We seem to be able to know certain claims like '$1+1=2$', 'if it is raining then it is raining' and 'every number has a successor' despite not having any non-question-begging argument for them. In contrast, it appears that we *would* need further argument to know other claims like Fermat's last theorem[1] or that smoking causes cancer. Call those claims not requiring further argumentation *acceptable starting points for a priori argument.*

Now, does this intuitive distinction between acceptable and unacceptable starting points for a priori reasoning track a principled difference in content, in virtue of which some propositions can be known without further justification and others cannot? Or does it ultimately reflect highly contingent facts about how human psychology relates to these propositions, e.g., what truths human beings tend to find obvious?

I expect the first option will strike most readers as prima facie more attractive than the alternative. We might be inclined to admit some minor exceptions to this claim to allow for beings intelligent enough to apply many of our inference rules in a single step, but on the whole it can seem obvious that the distinction between acceptable and unacceptable starting points isn't a personal or species relative notion. Instead we expect that this distinction could somehow be compellingly formulated without any appeal to our individual or species quirks.

[1]Fermat's last theorem says that no three positive integers a, b, and c can satisfy the equation $a^n + b^n = c^n$ for any integer value of n greater than two. It was proved by Andrew Wiles[Wiles, 1995] and we will assume, as is widely believed, that it is provable from Peano Arithmetic.

I propose to consider a pair of thought experiments that puts this expectation into doubt. I will then discuss a number of potential strategies to salvage a principled notion of acceptable starting points for a priori argument and the drawbacks to each of them. In light of these drawbacks, I will conclude by exploring the more radical option of allowing that the distinction between acceptable and unacceptable starting points partly reflects contingent facts about what humans find obvious. While initially unappealing, I will argue this later approach allows us to make sense of certain intuitive examples of a priori knowledge of contingent claims in situations which are very different from the classic cases discussed by Kripke[2]. It may even suggest a viable new method of attack on the old problem of external world skepticism.

## 1. A Thought Experiment

1.1. **Introducing the doctoroids and the math androids.** Claims which are not acceptable starting points can nonetheless feel immediately obvious at times. Think of what it feels like for an instant when you add $2 + 3$ and get 6, or of the generations of people who believed that two parallel lines in space necessarily intersected at most once. In view of this possibility, I want to propose two thought experiments involving creatures who find, not falsehoods, but additional truths immediately obvious.

In the first case the additional truths will be contingent empirical claims and in the second mathematical truths. While problems about justifying mathematical assumptions might be brushed off as a special case or explained away by relativizing the notion of proof, and empirical claims might be simply declared invalid starting points for a priori reasoning, taken together they clearly demand a more complex analysis.

[2][Kripke, 1972]

1.2. **Doctoroids.** Suppose that in order to save 6 years of medical education we engineer creatures who are disposed to find certain true propositions of organic chemistry and medicine immediately compelling, in phenomenologically much the same way the way that we find the claim $1 + 1 = 2$ immediately compelling. As soon as they consider the question of how smoking effects humans, they find it immediately obvious that it causes cancer. The doctoroids feel no compulsion to produce further justification for their judgements nor have they (generally) done or heard about the kind of empirical experiments which we would normally judge to be necessary to support these conclusions. Suppose further that we don't provide them with evidence that would allow them to deduce that we designed them to reach correct rather than incorrect conclusions.

One might worry that the doctoroids don't truly refer with the medical terms they use since, e.g., thinking about smoking requires certain causal interactions with lungs, tobacco or the like, which the doctoroids have not had. To avoid this objection let us further suppose that the doctoroids are raised like normal human children with all the exposure and linguistic mastery this requires, and they are only presented with the question about smoking and cancer once they correctly refer with those words[3].

[3]Alternatively, imagine that we start by replicating the brains of fully-functioning adult English speakers with amnesia about their past, and then modify these brains in such a way as to create the relevant dispositions to respond to english sentences and have feelings of obviousness. Perhaps initially, like Davidson's Swampman [Davidson, 1987], these creatures will not really count as thinking about smoking, cancer etc. For plausibly, thinking about smoking requires certain causal interactions with lungs, tobacco or the like, which the doctoroids have not had. Now, however, we switch the doctoroids on and allow them to take part in usual causal interactions with the world. Like Davidson's Swampman their talk will gradually pick up genuine reference as they interact with the world. In the course of this procedure they may also get some slight empirical evidence in favor of the medical propositions which they find obvious e.g. nothing in their experience seems to explicitly contradict the prediction of these theories. However, I take it that the amount of causal contact with the world which one needs to be able to *think about* the germ theory of disease or the kind of atomic theory embodied in the periodic table of the elements is far less than the experience needed to *justify* believing these claims. Thus we can choose some point at which the doctoroids have the minimal causal contact with the world needed for their words to have the usual reference but not enough experience

1.3. **Math Androids.** Along similar lines, imagine that in order to save time in graduate school, we engineer creatures who are disposed to find certain complex but truth preserving mathematical inferences immediately compelling, in phenomenologically much the same way the way that we find $1 + 1 = 2$ obvious. So for example, they find it immediately obvious and compelling to infer Fermat's last theorem (FLT) from any premises, in much the same way as we are willing to infer 'if it is raining then it is raining' from any premises or 'it is raining' from 'it is raining and it is snowing'. As in the case above, they don't ask for any further justification of these claims before accepting them. Indeed, they find it silly to search for the kind of proof that Andrew Wiles produced, since they can deduce the theorem in a single line. Just like the doctoroids, they too can't indirectly justify their inferences from evidence about our intentions in designing them.

1.4. **The Key Question.** Now that we've met the doctoroids and the math androids we can ask: do the doctoroids and math androids know the claims which they find immediately obvious? Prima facie there is something to be said for both options.

On the side of 'yes', the doctoroids' situation with respect to medical claims seems directly analogous to the situation that Descartes took us to be in with respect to mathematical and philosophical knowledge. According to Descartes, God, not being a deceiver, engineered us to find certain true propositions and valid inference methods obvious[4]. God was supposed to have put the dispositions to form true beliefs into us, 'like sparks in a tinder'. One might not buy Descartes' arguments that there really was such a benevolent creator. But it's pretty intuitive to think that, *if Descartes*

to justify the medical theories which they find immediately obvious. Now consider the epistemic status of the doctoroids' beliefs at this point.

[4]Descartes, Rene. Meditations on First Philosophy, translated by John Cottingham (Cambridge: Cambridge University Press, 1996). Book IV

*really was the beneficiary of this kind of divine benevolence*, he would count as having various kinds of a priori knowledge. However, (it seems that) the only thing which differentiates divinely-assisted-Descartes from the doctoroids and math androids is that he has cognitive faculties that reliably generate unthinking acceptance of true propositions about **fewer things**.

On the side of 'no', if we allow that the doctoroids can know claims like 'smoking causes cancer', this threatens to either trivialize or parochialize the distinction between a priori and a posteriori propositions. If the doctoroids know at all then they know a priori[5], for they lack any empirical evidence to justify the claim. If we accept the natural idea that all propositions which can be known a priori (by finite creatures like us) are a priori, this would entail that 'smoking causes cancer' is an a priori truth leaving us with a contradiction. Now we might try and wriggle out of this conclusion by insisting that 'smoking causes cancer' is a posteriori, despite the fact that the doctoroids can know it a priori. However, this leaves us with no way to distinguish a priori and a posteriori claims without directly appealing to what *we* happen to be able to know a priori[6].

An analogous problem exists for the math androids. Recall that the math androids accept FLT without any further argument, and believe any claims $\phi$ which they can prove from it. It seems plausible to say that[7] they know that $\phi$ in such cases if and only if the argument which they have is an adequate proof of $\phi$. But, intuitively, an argument for $\phi$ which takes Fermat's last theorem as an un-argued premise only counts as a proof of FLT$\supset \phi$ not a proof of $\phi$, and this kind of conditional proof is not adequate to confer knowledge that $\phi$. Thus we are pushed to say that the doctoroids do *not*

[5]To be precise they know that one can derive smoking causes cancer from no premises a priori.

[6]I presume that we don't want to jettison the idea that truths are a priori in virtue of being in *some* sense knowable by in an a priori fashion.

[7]assuming that they don't happen to posses any separate proof of $\phi$

know FLT and the claims which they have proved from it. One could, of course, try to count the math androids' proofs as acceptable because the conclusions follow necessarily from the premises, thus distinguishing the math androids from the doctoroids. However, this would require a radical revision indeed. Arguably taking a line like this on proof would remove the whole point of distinguishing between proofs and non-proofs, demanding a proof etc. since, by this criterion every necessary truth could be given a one line proof.

In view of these motivations to say that the doctoroids and math androids don't know, it is natural to look for some kind of principled explanation of why is it is epistemically acceptable for us to believe certain claims without further argument, which (unlike the simple Cartesian story sketched above) would preclude creatures like the doctoroids and the math androids from counting as knowing the claims which they find immediately obvious.

I will now consider a number of attempts to give such an account.

## 2. Attempts to draw a principled distinction

At least two different strategies for drawing a principled difference between our situation and that of the doctoroids and the math androids are possible. On the one hand, it might be that something directly about the content of intuitively acceptable starting points makes these truths epistemically acceptable to assume.

Alternately, it might be that we stand in a special relationship to the content of acceptable starting points, a relationship which it would be, in principle, impossible for the doctoroids or the math androids to stand in to the content which they find obvious. On this view, merely immediately assuming a proposition is never epistemically warranted, regardless of the content of that proposition. We are sometimes in a position to know claims like $1 + 1 = 2$ without further argument because we can arrive at these claims by deploying a special faculty, such as concept inspection or a special form of rational reflection. However, something about the nature of this process insures that it only works for certain kinds of truths. This gives us a principled distinction between propositions: the acceptable starting points for a priori arguments are exactly those truths which could in principle be delivered by deploying the relevant faculty.

I will discuss examples of both strategies in the subsections that follow.

2.1. **Infallible Method.** A first proposal is motivated by the traditional idea that good methods of a priori reasoning are ones such that it would be metaphysically impossible for them to lead one astray. Perhaps what allows us to know claims like $1 + 1 = 2$ without further argument is that it would be impossible for the method of assuming that $1 + 1 = 2$ to yield a false belief.

This proposal would succeed in explaining why, in my particular thought experiment above, the doctoroids don't know. However it fails to explain why the math androids don't know. Assuming Fermat's last theorem is just as infallible a method as assuming that $1+1 = 2$. More generally, the problem is that there are plenty of necessary truths which are not intuitively acceptable starting points claims. Truths like the right answer to the Goldbach conjecture[8], the right answer to the trolly problem[9] or the fact that that water is $H_20$ are also necessary truths, so it is just as impossible to form a false belief by believing these propositions as it is impossible to form a false belief by believing that $1 + 1 = 2$.

[8]The Goldbach conjecture says that every even number greater than 2 can be expressed as the sum of two primes.
[9][Thomson, 1976]

Note that I am only arguing here that considerations of infallibility and necessary truth can't suffice to account for the distinction between acceptable unacceptable starting points for a priori reasoning. This is not to deny that there may be some adequate principled account of the distinction between acceptable and unacceptable starting points which *implies* that all only necessary truths can be acceptable starting points. Indeed we will consider some options for such a principle in the sections below.

2.2. **Place in a Coherent Network of Obvious Feeling Beliefs.** A second natural idea is that what makes it epistemically acceptable for us to accept $1 + 1 = 2$ without further justification is that this belief forms part of a coherent network of other obvious feeling propositions within a largely empirically adequate picture of the world. Many obvious feeling mathematical claims not only feel obvious, but can also be derived from one another. For example, you can derive truths like $1 + 1 = 2$ from the Peano Axioms. And, more generally, there are various different ways of axiomatizing arithmetic which are inter-derivable. Thus it seems prima facie plausible that the fact that $1 + 1 = 2$ can be supported by other elements in a web of mathematical claims that feel obvious to us either heightens our justification for, or removes some kinds of defeater to, our knowing that $1 + 1 = 2$.

It is less obvious whether other intuitively acceptable starting points for reasoning (like 'I exist' or 'nothing is red and green all over') belong to similar webs of coherent intuitions, but perhaps they do. For our purposes, however, this doesn't matter.

We can already see that the criterion above cannot be used to support the idea that the intuitive distinction between acceptable and unacceptable starting points reflects a general constraint on what any finite creatures could know immediately. For, this criterion is too lax to exclude the doctoroids: the scientific theories which we upload into the doctoroids' brains are both coherent and empirically adequate. Just as in the mathematical case, one will be able to derive some of these medical beliefs from others and support generalizations with particular cases. Thus the doctoroids' medical beliefs will also form part of an internally coherent and a largely empirically adequate picture of the world.

2.3. **Result of a Faculty Whose Accuracy We Can Account For.** A third suggestion is inspired by the thought that what's intuitively bad about the doctoroids' situation is that they don't have any satisfactory account for how they could have gotten good intuitions about medicine. It was crucial in setting up the case above to specify that the doctoroids don't know that we formed their medical intuitions in response to reliable evidence. Thus, keeping in mind this contrast, one might propose that what makes propositions like $1+1=2$ acceptable starting points for a priori reasoning is that we can tell a plausible story which explains why we would have largely correct intuitions in this area.

This suggestion can be specified in two different ways, but neither specification of it is adequate.

If the proposal is that a person can only be justified in accepting claims that immediately feel obvious to them about a subject matter if they are now in possession of a good explanation for how they managed to get true beliefs about that subject matter, then this criterion is too strict. It's famously difficult to explain why our mathematical beliefs should have any correlation with their subject matter. Benacerraf points out in [Benacerraf, 1973] that realists about mathematics face a serious problem in accounting for why our beliefs about abstract objects like numbers and sets should have any correlation at all with the facts about these objects[10]. Presumably we don't want to say that no one can count as knowing that $1 + 1 = 2$ until they come up with an adequate response to the Benacerraf problem.

Alternately, one might propose that we are now justified in believing a priori truths if and only if we could, in principle, discover such a story. This way of spelling out the idea avoids over-strictness by allowing that we can be justified in believing that $1 + 1 = 2$ now, in virtue of the mere possibility of telling such a story. Unfortunately however, it also loses the distinction between ourselves and the doctoroids. For, plausibly, both we and the doctoroids could learn to account for our having accurate intuitions about mathematics and medicine respectively. We could learn that evolution shaped us to reason correctly about logic and arithmetic, and the doctoroids could learn that we shaped them to reason about medicine correctly.

One might point out that the doctoroids would need to get extra empirical evidence (which they don't now have in order) to arrive at an adequate account of how they acquired correct medical knowledge. However, one plausibly also needs empirical evidence to be justified in believing an evolutionary story about human psychology. And it seems perfectly possible that the right account of human knowledge of mathematics should require appeal to some such empirically discoverable process. But even on this hypothesis that evolution is needed to give a satisfying answer to the access problem, it would seem that these people who lacked adequate evidence for evolution could still count as having mathematical knowledge.

2.4. **Acceptability to All Thinkers.** A fourth proposal is to say that what makes some propositions acceptable starting points is that nothing

[10]One might be tempted to respond to this by giving up realism about mathematical objects but note that e.g. fictionalists face a similar problem in accounting for our ability to get true beliefs about what further truths $\psi$ must hold in the a fictional scenario characterized by some claim $\phi$.

could count as a thinker without being inclined to accept them. Here the idea is that the Doctoroids could "retreat" from their putative medical knowledge, and rationally evaluate whether they have such knowledge using mathematical and logical reasoning, but we cannot "retreat" from our logical intuitions to question whether they work. Perhaps what makes it warranted for us to believe certain claims (like logical truths) without having further argument for them, is the fact that these claims are so fundamental that if we retreated from that we wouldn't be reasoning at all. This suggests the following idea: one is justified in assuming a proposition without justification if and only if there is no possibility of retreating to a weaker theory in which to attempt to justify it.

The problem with this suggestion is that if any elements in our web of belief have the property of being 'impossible to retreat from' in the above sense, far too few of them do.

Firstly, it is far from clear that our sample claim $1 + 1 = 2$ has this property. Why couldn't someone accept, say, first-order logic without accepting any claims about mathematical objects and still count as a thinker?

Secondly, it is not even clear that the principles of classical logic are impossible to retreat from in the relevant sense. Maybe someone could accept intuitionistic logic but not the law of the excluded middle and still count as a thinker.

Finally, consider thinkers with sensory or conceptual impoverishment. Plausibly a blind person could lack the belief that 'nothing is red and green all over' and still count as a thinker. Yet 'nothing is red and green all over' is plausibly an acceptable starting point for reasoning. Similarly perhaps someone who has not yet learned some concept like ornate might be a thinker, yet they wouldn't believe any of the propositions we would express with claims about ornateness.

2.5. **Conceptual Necessity.** Inspired by the above objection, one might attempt to account for the distinction between acceptable and unacceptable starting points by limiting the above claim to thinkers who can count as thinking the relevant proposition. This gives us the suggestion that what makes some propositions acceptable starting points for a priori reasoning is that anyone who has the concepts needed to think about them would accept them.

Thus, for example, one might deal with the objections above by saying that no one can count as having the concepts of number or addition unless they also have certain beliefs including $1 + 1 = 2$.

Against this line of response, I first want to argue that it threatens the status of intuitively acceptable starting points like the claim that for all numbers $a$ and $b$, $a + b = b + a$. I take it that these claims *are* acceptable starting points for mathematical reasoning. Someone like Euler who accepted these claims without mathematical proof could count as knowing them and all the results which he proved from them.

On the other hand, it is not the case that anyone who could count as thinking about the numbers would have to accept that for all $a$ and $b$, $a+b = b + a$, any more than everyone who counts as thinking about the numbers would have to accept that there are infinitely many primes. For it would seem that someone who accepted the axioms of second order arithmetic and the Peano axioms could thereby count as thinking about the numbers. Yet if you only allow yourself to appeal to these axioms, the statement that for all $a$ and $b$, $a + b = b + a$ requires a proof at least as long as the standard proofs (in normal mathematics) that there are infinitely many primes.

More generally, the problem is that a number of *different* sets of mathematical premises and inferences seem intuitively sufficient to allow a person who accepts them to count as genuinely thinking about the numbers. Thus it is doubtful whether there any particular claim which must be believed by anyone who counts as thinking about the numbers. And it is almost certain that not all the claims which legitimate past mathematicians have taken as unarmed premises have this property.

We may be tempted to say that for all $a$ and $b$, $a + b = b + a$ is an acceptable starting point because, even though you can count as having the concept of number without believing this claim, anyone who counts as having the concept of number must accept *some* premises and inference methods from which this claim could be proved (though the relevant proof may be different for different people). However, if one loosens the requirement in this way, one loses the result that the math androids don't know Fermat's last theorem. For, as the existence of Wiles' proof demonstrates, Fermat's last theorem is provable from claims that are intuitively acceptable starting points. Thus, if this proposal can vindicate the acceptability of standard mathematical premises and inference proceedures by securing the fact that *they* are provable to anyone with the relevant concepts, it will follow that Fermat's last theorem is also provable to anyone with the relevant concepts - hence FLT will also be classed as an acceptable starting point by this criterion[11].

## 3. The Parochial Approach and Its Advantages

3.1. **Examples of the parochial approach.** Having shown how a range of natural attempts to make a principled content-based distinction between propositions that are and aren't acceptable starting points for a priori reasoning fail, I will now consider the alternative. Perhaps the intuitive class

[11]A final possibility along these lines would be to saying that acceptable starting points are ones that could be proved to anyone who counts as having the relevant concepts via a *short* argument. Thus for example, we might say that acceptable starting points in arithmetic include all statements that could be given a short proof to anyone who counts as having the concept of numbers. I think this is an interesting option, and faces some serious challenges, but I won't further discuss it here.

of acceptable starting points for a priori arguments *is* parochial in the sense that it reflects specific and contingent facts about human psychology, such as facts about which truths human beings are disposed to find obvious. There are a wide range of forms this kind of view can take so I will only sketch a few to give an idea what such theories might look like.

3.1.1. *Reliable Method Parochialism.* The first view I'd like to present amounts to a kind of a parochial twist on the reliable-methods idea introduced in section 2.1. We might say that any reliable process of a priori argument suffices to grant a subject knowledge. This view has the consequence that certain contingent truths can be known a priori. Assuming a contingent truth without further argument can be a reliable method of forming true beliefs, provided that that that relevant truth is sufficiently 'modally stable', in the sense that it holds at a suitably wide swath of the closest possible worlds to the actual world. Thus if the chemical and medical laws which the doctoroids assume are sufficiently basic, reliable etc. to hold at all close possible worlds, the doctoroids can count as knowing them a priori.

On this view, the notion of being an acceptable starting point for a priori argument functions something like the notion of edibility. What is edible for one person or group may not be for another, just as what is an acceptable starting point for a priori argument will differ from group to group. In particular, the doctoroids will be in a position to gain a priori knowledge of whatever can be proved from the (modally stable) chemical and medical laws which they find immediately compelling while (barring massive change in the psychology of what claims we find immediately obvious) we will not be able to come to know these propositions a priori. However, *the claim that the world is made of atoms will not be an a priori truth on this view* as we evaluate whether something is an a priori truth based on whether it is provable via the particular forms of reliable a priori argument which normal human beings tend to accept[12].

Unfortunately, this view has some slightly weird consequences. For one thing, it implies that someone who was lucky enough to be brutely compelled by methods of armchair reasoning which presumed the existence of a flying spaghetti monster in a world where that monster actually exists (and is suitably involved with the most fundamental laws of nature to exist at all close possible worlds) would count as knowing that the flying spaghetti monster exists.

3.1.2. *Necessarily Reliable Method Parochialism.* Alternately, one might attempt to avoid the weirdness of reliable method parochialism by insisting that the the relevant premises or inference methods must be not just reliable but infallible. This avoids the surprising and perhaps unattractive consequence of allowing the doctoroids to have a priori knowledge of contingent claims. However, it still allows the doctoroids to count as gaining knowledge from 'proofs' that take Fermat's last theorem as an unargued premise.

Also, note that taking this view raises an immediate question about why infallibility as opposed to mere reliability is needed for a priori knowledge. Unlike Descartes, most contemporary philosophers don't think infallibility is needed for knowledge in general. Nothing seems directly absurd about the idea that a merely reliable process of, e.g., sense perception should be able to generate knowledge. Thus, it can seem puzzling and unmotivated to simply say that a priori knowledge requires the use of an infallible faculty while empirical knowledge only requires a reliable one.

Many philosophers have, indeed, thought that justified a priori reasoning was infallible, and could deliver only necessary truths. But I take it that they thought this fact had a deeper explanation. Something about the particular

[12]holding fixed the facts about our own psychology

kind of faculty or relationship that delivered a priori knowledge (e.g., concept inspection, reflection on necessary conditions for all experience, framework stipulations) ensured that it wouldn't make sense for genuine uses of this faculty to ever deliver false results. They didn't just think that it was a brute fact that a priori knowledge required the use of infallible methods while using merely reliable methods could suffice to deliver empirical knowledge.

3.1.3. *Total Parochialism.* Lastly, one might just fix what we take to be acceptable starting points for a priori argument and insist that creatures have a priori knowledge of a claim only if their justification would be acceptable to us. This has the nice consequence of identifying a priori truths with claims that are knowable a priori. However, by fixing at this particular point in time what counts as an acceptable starting point, we run the risk that as we grow more intelligent we will expand that list but find ourselves unable to say we have a priori knowledge of the new conclusions we reach.

Here is what I mean. Suppose some group of people came to find more mathematical truths, or truth-preserving mathematical inferences, immediately compelling. By the criterion above they would not count as knowing these propositions, unless they had a proof of them from more basic mathematical claims which normal people find obvious. If made aware of this situation and persuaded of the truth of Total Parochialism about acceptable starting points, these mathematically privileged people would be forced to say things of the form, 'p but I don't know that p' in such cases. Yet as, e.g., Moore has argued this would seem to be a sentence which no one should ever assert.

Perhaps one can sweeten the pill by recognizing that people in this situation (people who, as it were, grew into becoming math androids) would promptly form a notion of 'knowledge prime' which behaved much like our notion of knowledge but allowed a slightly expanded range of true propositions to count as acceptable starting points.

3.2. **Attractions of the Parochial Approach.** While initially taking a parochial approach seems unappealing, it offers certain advantages. In particular, many ways of spelling out the parochial approach allow us to make sense of a priori knowledge for certain kinds of contingent truths. Consider the status of contingent claims like 'not PEASOUP'.

**PEASOUP** Everything outside a 5 foot radius around you is pea soup, but when you walk by the soup forms up into ordinary material objects around you, making it appear as if the normal laws of physics govern the entire world.

We seem to know that we aren't in an "induction bubble", where the laws of physics behave one way within a 5 foot radius of us, but then radically change outside that radius in a way that's systematically undetectable, as per PEASOUP. But this (apparent) knowledge does not seem to depend on any experience. Thus we seem to know some contingent facts a priori, and it is an attractive feature of the above deflationary accounts that they make room for such knowledge[13].

This kind of apparent a priori knowledge of contingent truths contrasts with Kripke's familiar examples of contingent a priori knowledge that e,g. the meter stick in Paris is a meter long[14], insofar as there is nothing intuitively 'linguistic' about the contingent propositions being known. As such, it can seem hard to account for. However, both the first and third versions

[13]Hawthorne defends the idea of "deeply contingent" knowledge of a priori propositions against certain objections in [Hawthorne, 2002]. Interestingly [Turri, 2011] responds that Hawthorne's defense fails to address the intuitive objection that one needs further "evidence or argument" to know deeply contingent claims a priori, and one lacks such evidence or argument in Hawthorne's cases. Thus Turri's objection turns on the assumption that deeply contingent truths cannot be acceptable starting points for a priori argument.

[14][Kripke, 1972]

of parochialism specified above directly imply that commonsensical people who find not PEASOUP obvious can indeed thereby count as knowing that not PEASOUP without further argument. I take it this is an attractive feature of such views.

However, along with this advantage this approach requires one accept a certain weirdness. In particular, we can know skeptical scenarios are false even though we are unable to produce any argument or consideration which the skeptic will find convincing as to why they are false. We know they are false based simply on brute intuitions. Unease at this strangeness may be partially alleviated by noting that a parochial approach undercuts the skeptical challenge itself. After all any non-Pyrrhic skepticism must base its argument on distinguishing the acceptable starting points which the skeptic herself appeals to from those potential starting points which simply deny the skeptical position. For example, the external world skeptic owes us an account of why it is epistemically justified for them to assume whatever philosophical claims about knowledge, conceivability etc. they need to get their argument off the ground, but not for us to assume that we are not brains in vats. If no such principled distinction can be drawn then our inability to answer this ill-formed challenge seems less troublesome.

## 4. Conclusion

In this paper, I have considered a number of different attempts to sustain the idea that that the intuitive distinction between acceptable and unacceptable starting points for a priori arguments tracks a principled general constraint on what could be known a priori by all finite creatures. I have argued that all these strategies face serious problems, and indeed each one but the last fails.

I have also outlined an alternative deflationary approach on which the distinction between acceptable and unacceptable starting points for arguments turns out to reflect deeply *un*principled contingencies of human psychology, and shown that this approach has some attractions.

## References

[Benacerraf, 1973] Benacerraf, P. (1973). Mathematical truth. *Journal of Philosophy*, 70:661–80.

[Davidson, 1987] Davidson, D. (1987). Knowing one's own mind. *Proceedings and Addresses of the American Philosophical Association*.

[Hawthorne, 2002] Hawthorne, J. (2002). Deeply contingent a priori knowledge. *Philosophy and Phenomenological Research*, 65(2):247–269.

[Kripke, 1972] Kripke, S. (1972). *Naming and Necessity*. Blackwell.

[Thomson, 1976] Thomson, J. J. (1976). Killing, letting die, and the trolley problem. *The Monist*, 59:204–217.

[Turri, 2011] Turri, J. (2011). Contingent a priori knowledge. *Philosophy and Phenomenological Research*, 83(2).

[Wiles, 1995] Wiles, A. (1995). Modular elliptic curves and Fermat's last theorem. *Annals of Mathematics*, 141:443–551.