
What Would a BIV Do Differently? *A Pragmatist Defense of Contextualist Fallibilism*

ABTIN AMINZADEH DEZFULI

In "Solving the Skeptical Problem", Keith DeRose offers a contextualist response to a possible formulation of the skeptical hypothesis about knowledge. I will here outline his position in order to demonstrate the potential in the contextualist approach to effectively solve the skeptical puzzle. I will, however, go on to argue that the contextualist response as formulated by DeRose falls short of achieving its persistently elusive goal. In this, I will follow David Lewis, in "Elusive Knowledge", in order to explain how the type of contextualist solution offered by DeRose is inherently self-defeating. I will then suggest the introduction of a pragmatist understanding of knowledge into the contextualist picture. Shifting towards fallibilism, I will argue that in light of pragmatist considerations, the skeptical puzzle loses much, if not all, of its threatening significance.

DeRose is justified in his initial observation that any satisfactory solution to the skeptical argument should provide an account of the undeniable plausibility of its premises and the entailed conclusion. But why should an admittedly plausible argument be so troubling in the first place? The skeptic is, after all, merely inviting us to concede that we could possibly be deceived by the senses. On the face of it, we should be able to concede this point without facing a foundational crisis about knowledge. The problem, ironically enough, arises with Descartes' convincing argument to the effect that entertaining the notion of deception leads to the possibility of overall 'systematic' deception. Descartes' evil demon, or its contemporary counterpart, the brain-in-a-vat (BIV), are scenarios about precisely this type of situation: the case where all our beliefs about the world are, in one way or the other, potentially false. It is in this sense that the skeptical hypothesis, hereafter SH, remains a lark in the side of modern epistemology.

Consider DeRose's formulation¹ of SH (where O is an ordinary knowledge claim such as "I have a hand", and H is a skeptical proposition like "I am a BIV"):

P₁: I don't know that not-H.

P₂: If I don't know that not-H, then I don't know that O.

C: Therefore, I don't know that O.

DeRose correctly observes that SH, formulated as such, is at least at first glance a sound argument with intuitively plausible premises. However, it stands in direct conflict with an equally plausible knowledge claim that P₀: I know that O (e.g. I know that I have a hand). The problem, therefore, is that we are dealing with a *very* good argument with a *very* unacceptable conclusion. Setting aside the Moorean reaction of disputing the greater or equal plausibility of SH compared to P₀, the difficulty of the task at hand lies in the fact that the usual method of attacking the validity of premises or the soundness of the argument — namely, by denying any of P₁, P₂ or C — is an inherently implausible move in itself. The attractiveness of DeRose's contextualist solution, as I will sketch out below, is located in the way in which it manages to avoid taking such an unappealing stance.

DeRose's contextualist solution involves two major steps. The first is to investigate the reasons for the intuitive plausibility of both P₁ and P₂, and the second to demonstrate how we can accommodate that plausibility while keeping intact our belief in P₀. To this end, DeRose relies heavily on two central concepts: the Subjunctive Conditional Account (SCA) and the Strength of Epistemic Position. The SCA, taken from Nozick's "Philosophical Explanations", effectively argues that in order for a belief to count as knowledge, it should be 'sensitive' to the condition of being false.² In other words, my true belief that P is not my knowledge that P if it were the case that I would maintain that P even if not-P. Knowledge should track the truth in a specific way: it should be sensitive to its possible loss.

1. Keith DeRose, "Solving the Skeptical Problem," in *Skepticism: A Contemporary Reader*, ed. Keith DeRose and Ted A. Warfield (New York: Oxford University Press, 1999), 183.

2. Robert Nozick, "Philosophical Explanations," in *Skepticism: A Contemporary Reader*, ed. Keith DeRose and Ted A. Warfield (Oxford: Oxford University Press, 2005), 156-182.

SCA thus explains the plausibility of P_1 : our belief that not-H does not track the truth. The SCA by itself, however, would lead us to accepting logical absurdities. If we apply SCA to P_2 , we could argue that the conditional is broken because, despite the truth of the antecedent, the consequence is insensitive and as such false. We are therefore left with an abominable conjunction: the antecedent of the conditional and P_0 . Furthermore, we would have violated our initial promise to allow for the plausibility (and not the truth) of both P_1 and P_2 .

To avoid all this, DeRose introduces a contextualist condition to his argument for the plausibility of P_1 . It is only in a context with a heightened standard of knowledge that P_1 is actually true. In an ordinary context, that is, almost anywhere outside a seminar on epistemology, the context-determined strength of the conversers' epistemic position doesn't require P_1 to be true. This is the trick the skeptic has been playing on us all along: a change in the relative strength of the epistemic position caused by introducing the SH to the conversation. The same line of thought can be extended to P_2 . Given the newly raised epistemic standards, the conditional is indeed true. The contextualist explanation, therefore, is that the skeptic's introduction of H to the conversation has raised the epistemic standards to an extraordinary position where our belief that not-H is now insensitive. However, DeRose is quick to point out that the skeptic has by the same token raised the epistemic standard of knowing P_0 so high that we no longer can claim to know it.

Hence, DeRose manages to avoid endorsing the abominable conjunction. But has he solved the skeptical puzzle? So far it might have seemed that all he has done is to argue for P_1 and P_2 . Yet it should be noticed that the conclusion of the SH is now only to be followed under heightened epistemic standards. As such, the sense of the word 'know' in C is such that it poses no direct threat to our simultaneously maintaining that P_0 . DeRose thus argues that the combination of SCA and the notion of epistemic standards has successfully solved the skeptical problem. It has only to be realized that the skeptical worries are only warranted in an extraordinary truth insensitive context and not that of meaningful everyday conversation. In other words, the SH can be raised in any given context because it is context-insensitive in an important sense: it is

only plausible in a context that precludes our ordinary notion of knowledge.

As David Lewis points out, however, this solution places epistemology itself in a rather awkward position. The entire strength of the contextualist view rests on the basic observation that some possible alternatives — specifically those related to insensitive, systematic deception — are properly ignored in ordinary conversations. If I see what appear to be zebras in a zoo, I am justified in ignoring as irrelevant the alternative possibility that I am actually seeing cleverly painted mules, placed there by inexplicably deceptive zoo authorities. This seems to be a reasonable position insofar as the skeptical alternative has not been raised. As such, Lewis claims that ordinary knowledge of this type is infallible by virtue of having ruled out properly ignorable alternatives. It is crucially different from fallibilism, since in order to falsify the claim that S knows that P we must change the context of knowledge by heightening the epistemic standards. Indeed, I do not know that those animals are zebras if you mention the possibility of a paint job, since this alternative is no longer properly ignored. However, when I said they were zebras two minutes ago, I did know, and know infallibly, that they were zebras as opposed to antelopes.

But how do we know which alternatives are properly ignored in the first place? Here, DeRose is of no help. Lewis, however, makes a valiant effort to give us a rough idea of what some rules of proper elimination of possibilities will look like. The rules include criteria for actuality, belief, resemblance, reliability, method, conservatism, and attention. I believe none of the rules described by Lewis will give him the result that he needs for his infallibilism, that is, the possibility of properly ignoring the skeptical alternatives in all ordinary conversations. For the purposes of the objection I raise below, however, it will suffice here to consider the ‘rule of attention’. As Lewis states it, the rule specifies that “when we say a possibility is properly ignored, we mean exactly that; we do not mean that it *could have been* properly ignored.”³ This entails, of course, that the so-called far-fetched possibilities are relevant to the

3. Keith DeRose, “Solving the Skeptical Problem,” in *Skepticism: A Contemporary Reader*, ed. Keith DeRose and Ted A. Warfield (New York: Oxford University Press, 2005), 230.

context if they are for whatever reason not ignored by the conversers. To preserve his infallibilist view, Lewis suggests that *we* can nevertheless properly ignore the skeptical possibilities when discussing the knowledge claims in a conversation, *even if they are not ignored by the interlocutors themselves*. We are, after all, conducting a different conversation, in a different context, focusing attention on different alternatives of our own choosing.

I will allow Lewis himself to ask the immediately obvious question: “Don’t you smell a rat? Haven’t I, by my own lights, been saying what cannot be said?”⁴ The answer is yes on both counts. When discussing epistemology, we are discussing exactly the kinds of alternatives that we claim are being ignored in ordinary conversations. Lewis admits this self-destructing quality of the epistemological inquiry, yet maintains that despite this, his arguments remain correct. How so? Lewis believes that he has only momentarily “bent the rules” in talking about the contextualist position and ignoring the alternatives when in fact they were lurking in the background of the discussion. When we return to ordinary conversation, we can easily switch back to ‘properly’ ignoring the skeptical alternatives and, lo and behold, the skeptic withers away.

But is this a satisfactory response to the self-deconstruction charge? I think not. The crucial presupposition fueling the contextualist’s confidence in dismissing this charge is the global irrelevance of the skeptical alternatives. The possibility that I am seriously deceived will *never* be relevant in the context of an ordinary conversation and as such, our P_os are placed firmly out of the skeptic’s reach. This is what Lewis has in mind when arguing that knowledge is by definition infallible. As such, the burden remains on the contextualist to demonstrate the *absolute* context-independence of the skeptical alternatives. The contextualist, however, insofar as her position has any merit, is unable to move in this direction. Such an argument would commit her to the absurdity that she has all the while been talking outside any given context. For if the skeptical alternatives do make sense in some context — as they

4. David Lewis, “Elusive Knowledge,” in *Skepticism: A Contemporary Reader*, ed. Keith DeRose and Ted A. Warfield (New York: Oxford University Press, 2005), 237.

do, since the contextualist is discussing them — then it is indeed a peculiar move on the part of the contextualist to denounce them as absolutely context-independent. The *sotto voce* proviso would not suffice, since however soft the voice, something is being said, thus creating a context that cannot be essentially severed from ordinary life. In doing this, the epistemologist would be ascribing other-worldly qualities to philosophical discourse, which would imply that no one outside the seminar would ever understand the SH. But, then, how do philosophers get into this isolated island in the first place?

So, where does this last move leave us? We cannot simply capitulate by accepting C, since we will thus have made no progress in dealing with the puzzle that we started with, viz., the conflict between P_0 and C. Besides, we have already seen much promise in the contextualist solution. We have come too close to give up so easily. Often when a philosophical argument has fallen short, all it needs is a pragmatist boost. I will here try to give that boost to the contextualist solution.

C.S. Peirce, in “How to Make Our Ideas Clear”, argues that the content of our beliefs are a momentary appeasement of doubt and a resultant disposition to act. To say that I believe that P means that I am prepared to act in a manner compatible with the fact that P. Our consideration about the justification of a belief, therefore, includes the rule, “Consider what effects, that might conceivably have practical bearings, we conceive the object of our conception to have.”⁵ Peirce then goes on to suggest a scientific method of investigating these practical bearings, one that would ultimately lead to truth:

Different minds may set out with the most antagonistic views, but the progress of investigation carries them by a force outside of themselves to one and the same conclusion. This activity of thought by which we are carried, not where we wish, but to a fore-ordained goal, is like the operation of destiny. No modification of the point of view taken, no selection of other facts for study, no natural bent of mind even, can enable a man to escape the predestinate opinion. This great hope is embodied in the conception of truth and reality. The opinion which is fated to be ultimately agreed to by all who investigate, is what we mean

5. C.S. Peirce, “How to Make Our Ideas Clear,” *Popular Science Monthly* 12 (January 1878): 290.

by the truth, and the object represented in this opinion is the real. That is the way I would explain reality.⁶

How would this conception of belief and truth help us in epistemology? It helps by posing a possible contextualist explanation for our holding on to P_o , while accepting the relevance of SH to ordinary conversations (which was so devastating to the contextualist solution as offered by DeRose et al). Consider the case of driving in a facsimile barn land. DeRose maintains that despite the possibility of facsimiles my knowledge that I am driving through real barns attains the status of infallible knowledge since the facsimile alternative is properly ignored. Oh, unless of course, it is not ignored. The pragmatist, however, can readily accept the fallibility of the knowledge that she is looking at real barns, but add that *in the context of driving through a barn-land, the practical consequences of the barns being real or fake are effectively the same*. Now, if I were to settle down in this area, I would need to make further inquiries, which may or may not ultimately show that the barns were fake. As I was driving through, however, perhaps with an epistemologist friend of mine who was discussing the very same scenario with me, the possibility of deception had no practical bearing. In other words, for pragmatic purposes — and is that not all that matters anyways? — I knew that I was driving through barns. Having already incorporated the contextualist rules, my belief that P is a piece of fallible knowledge that P if and only if in the present context the possibility that not-P bears no immediate practical consequences for me.

Notice that the pragmatic solution stays within the defined limits of not attacking the plausibility of either of the two premises. Endorsing the fallibilistic characteristic of knowledge allows for this consistency. The pragmatic turn simply introduces a different way in which context defines the sense of the word 'know', even if the epistemic standards are raised to include far-fetched possibilities of deception. The pragmatist is asking the skeptic a legitimate question: "So what do you suggest we do differently, considering the possibility you just raised?"

6. C.S. Peirce, "Make Our Ideas Clear," 292.

As such, the pragmatist can happily embrace fallibilism while dismissing the immediate skeptical “Aha, I got you now” claim that she knows nothing because she does not know which of her beliefs actually constitute knowledge. Whether a belief constitutes knowledge is determined by the practical dynamics of the context, which makes justification an essentially futuristic project. The question of whether I am a BIV is *usually* quite detached from possible practical implications, and as such, properly dismissed. Yet, if the context was such that it would actually bear practical consequences, then the BIV hypothesis should be put to scientific investigation. At the ideal end of inquiry it may be proved or disproved that we are BIVs, maybe because Neo would manage to explode the Matrix, or the demon would get bored and decide to let us know of its existence for the sake of variety. As already conceded, our knowledge in this sense is infallibly fallible.

At the present time however, it seems that the BIV hypothesis is — as all claims about overall and systematic deception — designed specifically to elude empirical investigation. As such, the skeptic is being too clever for her own good. DeRose’s bold skeptic is thus defeated, not because she is making a statement about systematic falsehood that is unappealing to the anti-realists, but simply because her scheme bears no practical consequences for the knower. Since what we believe has relevant content only insofar as it determines the way we act, the skeptic is left with the burden of demonstrating the practical bearing of her particular hypothesis. If she manages to manipulate the context enough to do so, then we would remind her of our fallibilism in the first place and ask her to join us in the scientific inquiry she has shown to be pragmatically viable.

The skeptic has therefore rendered us a great service. By introducing the possibility of systematic deception, she has made us realize a serious flaw in the platonic picture of knowledge and truth. The demand for infallible knowledge should be given up, but given up in such a way that would capture the practical and evolutionary significance of the term ‘knowledge’. Pragmatic fallibilism provides this option.