

Contextualism and closure

N. D. Smith

Lewis and Clark College,
0615 SW Palatine Hill Rd. Portland, Oregon 97219, USA

For citation: Smith N. D. Contextualism and closure. *Vestnik of Saint Petersburg University. Philosophy and Conflict Studies*, 2020, vol. 36, issue 2, pp. 291–299. <https://doi.org/10.21638/spbu17.2020.207>

Skepticism confronts us with a paradox (sometimes known as “the skeptical trilemma”), a version of which follows: (1) I know that I am working on a computer right now; (2) I know that knowing that I am working on a computer right now logically implies that I am not being deceived or manipulated in the way that skeptical hypotheses imagine. (This implication is called “closure under known logical implication”); (3) I do not or cannot know that I am not being deceived or manipulated in the way skeptical hypotheses imagine. The paradox of skepticism is that these three statements are logically incompatible. A relatively new movement in epistemology called contextualism proposes that we can accept all three of the claims in the trilemma, by recognizing that they are not all true within the same epistemic context. Briefly, contextualists claim that we can know in ordinary contexts, but cannot know that we are not being deceived or manipulated in a skeptical scenario, but the latter fact is true in a different epistemic context than the ordinary knowledge that we might have. Closure under known logical implication will remain true, but only insofar as the implications involved are alternatives that belong to the same epistemic context as the original knowledge claim. In this paper, I claim that contextualism’s account of how epistemological contexts change, together with its acceptance of closure, is implausible.

Keywords: epistemology, knowledge, closure, contextualism, fallibilism, logical implication, skepticism.

The paradox of skepticism

It has now become something of a commonplace to note that global versions of skepticism work by generating scenarios (skeptical hypotheses) that seem to show that we cannot know what we take ourselves to know. For example, a skeptic might remind us of the evil demon hypothesized in Descartes’ first *Meditation*, and argue that we cannot know if we are either deceived by such a demon (in which case what we believed would fail the truth condition for knowledge), or even manipulated by such a demon (which would defeat our justification for our belief). But the skeptic goes on to argue that we cannot know even if neither of these hypotheses are false, so long as we are not in a position to know that they are false. The skeptic manages this aspect of their view by appealing to what has come to be known as the closure principle¹.

(Closure) If S knows q and S knows that q entails not-h, then S knows not-h.

The assumption of closure is vital to the skeptical argument, because the way skeptical arguments work is to fill in an example of an entailment, from some candidate for knowl-

¹ The following formulation comes directly from [1, p. 93–94].

edge, which apparently *can't* be known. So, let “q” be the proposition that S has a hand in front of his face, and let “h” be the skeptical hypothesis that S is deceived or manipulated by Descartes’ evil demon. According to closure under known logical entailment, if some epistemic agent (S) actually knows that he has a hand in front of his face, and knows that having such knowledge entails that he is not deceived or manipulated by an evil demon, then he can simply work to the conclusion, via *modus ponens*, that he knows that he is not deceived or manipulated by an evil demon. That is, in the following argument, one can work from premise (1) to the conclusion (3), using the closure principle as premise (2):

- (1) S knows that q.
- (2) S knows that q entails not-h.

Therefore,

- (3) S knows that not-h.

Indeed, many fallibilist internalists have argued for something like this as an answer to skepticism². But skeptics claim that the conclusion is plainly false: no one can really *know* that some global skeptical hypothesis is *false*. So, they use this same argument, negating the conclusion (3) and work a *modus tollens*, instead: given the negation of (3) and the truth of closure (2), they can validly derive the skeptical conclusion that (1) is false. It would appear that most if not all of what we ordinarily think we know can be a substitution instance for q in the above argument, which is why skeptical arguments seem to have such a powerful intuitive appeal.

The contextualist approach to the challenge of skepticism has received a great deal of attention in the past decade or so³. Contextualism, roughly, is the claim that the standards by which we are supposed to judge claims of knowledge vary according to different epistemic contexts, in which different epistemic considerations might apply. The factors that can create different contexts and thus different standards of evaluation of knowledge claims, might include, for example, how much risk is engaged if a claim of knowledge should turn out to be false. An example of this sort of case is given by Stewart Cohen:

Mary and John are at the L. A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and responds, “Yes I know — it does stop in Chicago”. It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, “How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute”. Mary and John agree that Smith doesn’t really know that the plane will stop in Chicago. They decide to check with the airline agent [9, p. 58].

What has happened in this case, according to Cohen, is that in ordinary epistemic contexts, the fact that Smith actually checked his itinerary would be enough for us to grant his knowledge claim. But Mary’s concerns, together with the practical importance to Mary and John that they make sure to get this one right, lead them to regard Smith as not know-

² This approach is often associated with G. E. Moore. For a contemporary version of this sort of fallibilist inference, see [2, p. 133–134].

³ See, for instance, [1; 3–8].

ing. The epistemic context, Cohen claims, has shifted, resulting in a requirement of higher standards for them to accept a claim as knowledge.

The appeal of contextualism is that it claims to be able to allow all of the steps of the earlier argument we considered, while not authorizing either the skeptic's rejection of premise (1) or the easy *modus ponens* first described, which could allow someone to affirm conclusion (3) with no better evidence against some skeptical hypothesis than ordinary experience would supply. That is, according to contextualism, premise (1) will be true in epistemic contexts, whereas premise (3) will *never* be true in ordinary skeptical contexts. Hence, the closure principle must be keyed to the relevant epistemic context: some considerations will be relevant to given epistemic contexts, and closure under known logical entailment will only operate when the alternative considered (substitution instance for *h* in the above schema) is relevant to that context. In brief, if some alternative possibility (*h*) is one that someone (*S*) cannot know or at least does not know is false in that context, then neither can the entailment of its falsehood be known by *S* in that context. The conclusion we should draw, according to contextualists, is that in any such case, the introduction of that alternative possibility changes the context into one with higher epistemic standards for knowledge. The introduction of skeptical hypotheses is thus an example of a context-changer. Contextualists argue that recognizing this about how alternative possibilities can change contexts explains why we find it so plausible to think that we do actually know many things (in ordinary contexts, for example), but also cannot know that some skeptical hypothesis is false. Closure, they claim, is maintained at every epistemic level, with the proviso that the only alternatives one can rule out are those that are relevant to that context. In a context in which consideration of some skeptical hypothesis is engaged, closure shows why the fact that we cannot know the scenario to be false also allows skeptics (rightly, in that context) to reject a knowledge claim we might otherwise have granted (that is, in a different context).

Now, non-contextualists have generally attempted two different sorts of replies to skeptical arguments. Some anti-skeptics (especially those inclined to positions widely known as “externalist” or “naturalist” accounts of knowledge) have tried to argue against the application of the closure principle. A simple example of this approach can serve to make clear roughly how it works⁴. Reliabilism contends that knowledge is true belief that is generated or sustained by reliable cognitive processes — that is, by cognitive processes that reliably generate or sustain true belief. If *S* forms the belief that he has a hand in front of his face (*q*) on the basis of ordinary perception, and ordinary perception reliably generates or sustains true belief (which it will for any epistemic agents who are not actually in some skeptical scenario), then we can conclude that

(1) *S* knows *q* is true.

Now it may be that *S* enjoys no reliable cognitive process that can assure *S* that he is not in some skeptical scenario. Hence, it will also be true that

(3) It is not the case that *S* knows that not-*h*.

⁴ Indeed, the example that follows is not merely a simple one, but is, indeed, a simplistic one. I do not intend herein to survey all of the approaches to skepticism and closure, however, but wish instead to focus on the contextualist approach. I offer simplistic samples of other approaches only to motivate the contextualist response to them.

But surely it cannot both be true that *q* and that *h*—it cannot both be true that *S* has a hand in front of his face and that *S* is a brain in a vat. And we can well imagine that *S* is quite aware of this incompatibility, as per (2). If we deny closure, however, this incompatibility does not yield skepticism. So, those who deny closure assert that *S* can know that there is a hand in front of his face, *even though* he does not know that he is not in some skeptical scenario and also *S* knows that he cannot know that there is a hand in front of his face *if he is* in a skeptical scenario. In essence, the position endorsed by those who deny that knowledge presumes closure is this: We can know *q* even though we do not know everything that we know follows from the truth of *q*.

Various clever and at least somewhat persuasive arguments have been given for other replies to skepticism, but contextualists have nonetheless resisted all such arguments. The contextualist approach against both sorts of replies to skepticism is essentially the same: The contextualist reminds us of just how intuitively compelling and troubling the skeptical argument is⁵. If one of the premises of the skeptical argument is simply false, the strong appeal of such arguments would be inexplicable. Against those who deny closure, contextualists have reminded us just how very intuitive the idea of closure is: If we really do know and recognize that something (not-*h*) is a *logical entailment* of some proposition that is a candidate for knowledge (*q*), how can we really claim to know the knowledge-candidate (*q*) if we grant that we do not and (given our circumstances or epistemic conditions) *cannot* know what is entailed by the knowledge candidate (not-*h*)? The idea of closure under *known* entailment, we should recall, is so intuitive that generations of philosophers have been troubled by the challenge of skepticism. Against Moorean fallibilists, too, contextualists have reminded us of the fact that skeptical scenarios invariably do challenge and trouble us. Whereas no one in good mental health really *worries* that he or she might actually be deceived or manipulated by an evil demon, the very idea that we might be invincibly deceived is not simply incredible on its face. Indeed, even if we are not at all inclined to believe that we are invincibly deceived, the limitations of our cognitive systems seem obvious enough to anyone who thinks at all about them, and so the idea that such limitations might be far greater and far more significant than we suppose does not really seem to be something we can simply dismiss, on the ground that we are also convinced we know many things.

One advantage of the contextualist approach, then, is that it allows us to affirm each of the claims that lead to the skeptical paradox. The contextualist provides an account that preserves the intuitive appeal of the closure principle, the skeptic's insistence that we do not and cannot know that we are not invincibly deceived, and yet does not concede the skeptic's conclusion that we do not know many things. But despite the appeal of this view, I shall now argue, two features of the contextualist approach cannot be reconciled, and one of these two features is the contextualist's acceptance of the closure principle. The other, which I must now explore in a little more detail, is the contextualist's way of understanding epistemic contexts.

⁵ So, as Cohen says: "The burden of the fallibilist is to resolve these puzzles and paradoxes in a way that preserves the truth of our everyday knowledge attributions. But a satisfying resolution requires an explanation of why the paradox arises—an explanation of why we have the intuitions that saddles us with the paradox. As I noted, in the case of the skeptical paradox, it is not enough to simply conjoin (3) with a second proposition of the inconsistent triad and then infer the denial of the third. Such a "resolution" would not tell us how the paradox arises in the first place" [1, p. 94].

Contextualizing epistemic standards

As is now well known, the way in which contextualists reply to skepticism is to point to the way in which consideration of skeptical scenarios seems to raise the presumptive standards for knowledge beyond what normal human beings can meet. Again, very simply (or simplistically), when we do attribute knowledge to someone (S) — for example, the knowledge that he has a hand in front of his face (q) — we do not seem to be considering the possibility that we or that S might actually turn out to be brains in vats. The context of what might be called “ordinary” knowledge claims, let us say, is the “ordinary” context. We should therefore understand the claim that S knows that q as one that is *made* within a certain context of cognitive discourse⁶.

But contextualists argue that the consideration of skeptical scenarios puts cognizers into a *different* epistemic context, within which different standards for knowledge apply. When asked to consider if we might be invincibly deceived, the relevant epistemic standards seem to require certainty, yet we find that we cannot be completely certain that we are not invincibly deceived. Because ordinary contexts do not require us to consider skeptical scenarios, and because the standards for disqualifying skeptical scenarios seem to be so high that we are inclined to think we do not satisfy them, we can be confident, according to the contextualist, that the ordinary epistemic context and the epistemic context in which we consider skeptical scenarios are different ones. Hence, although the contextualist will insist that closure under known logical inference applies within each epistemic context, it does not hold *across* epistemic contexts. And this is how the contextualist replies to the skeptic. In this way, what appears to be the inconsistent set of propositions we wish to affirm, which constitute the problem of skepticism, they never actually form an inconsistent set of judgments — for they are only true within specific epistemic contexts, and we can remain consistent by denying one or more of them within each of the epistemic contexts that apply to them. This is a nice result for contextualists, because it allows them to claim that we know lots of things, while also allowing them to recognize the power of skeptical scenarios to move us to deny knowledge.

A counterexample to contextualism in ordinary epistemic contexts

Contextualists seem to have a plausible response to the extraordinary standards the skeptical scenarios appear to require of us. But closure under known entailment remains a problem for the theory.

On a recent trip, I brought along my laptop to get some work done. I noticed, however, that my hotel room had not been supplied with a towel, so I decided to go down to the lobby to let the people at the registration desk know that I needed a towel. The following set of claims relates to the short time I was out of my room:

(Closure) If Smith knows that his laptop is in his hotel room, and Smith knows that his laptop's being in his hotel room entails that the laptop has not been stolen since he left his hotel room, then Smith knows that no one has stolen his laptop since Smith left his hotel room.

⁶ It is the context of the *subject* of the attribution (S), rather than that of the attributor, that matters. See [10].

- (1') Smith knows that his laptop's being in his hotel room entails that the laptop has not been stolen since Smith left his hotel room.
- (2') Smith does not know that no one has stolen his laptop since Smith left his hotel room.
- (3') Smith knows that his laptop is in his hotel room.

The contextualist grants that *any* knowledge claim requires warrant, and because they resist externalist accounts of warrant (because such accounts reject closure under known entailment), contextualists seem generally to regard warrant as consisting in some appropriate or sufficient justification. In the case of 3', the contextualist would grant my claim of knowledge, as my justification for the claim is certainly more than adequate for ordinary contexts—not only do I actually *have* substantial justification for this claim, since I had only just left the room minutes ago, and remembered very vividly exactly where in the room I left the laptop, and so on. So, it is plain that a contextualist would accept my claim in 3', at least within ordinary epistemic contexts.

But there are other things that I also know this case, including the banality that my laptop will not now be in the room if someone has taken it from the room since I left it there. Unlike skeptical scenarios, which contextualists argue considerably raise the stakes and epistemic standards customary and applicable to ordinary epistemic contexts, it seems plain that the observation that sometimes laptops are stolen does not require any elevation of epistemic standards, or for that matter, any significant consideration of anything extraordinary. Moreover, recognition of the obvious observation that when something is stolen, it will not remain in the place where it was before it was stolen, hardly qualifies as requiring epistemological sophistication. As a result, I contend that the knowledge of the entailment given in (1') must also qualify as knowledge the contextualist must regard as applicable within ordinary epistemic contexts.

The problem for contextualists, I take it, is now obvious, for by (1'), (2'), and (Closure), I should be able to derive the consequence that I do, indeed, know that my laptop has not been stolen since I left my hotel room. But it does not seem that I do—how could I know that, since I have not been in my hotel room for several minutes, and it does not take long to steal a laptop. How do I know that some thief has not gained access in my absence and made off with my laptop?

If I am right that the contextualist is committed to granting claims of knowledge such as (3'), and I am also right that *even in ordinary epistemic contexts* no one knows that possessions not *right now* under one's watch have not been stolen, then the contextualist's commitment to closure under known entailment generates the very problem of inconsistency—within a single epistemic context, in this case, what contextualists call the ordinary epistemic context—that his or her reply to skepticism was supposed to avoid.

A survey of the options and conclusion

Of course, contextualists could respond to this objection by abandoning the commitment to closure under known entailment. But the cost of such an abandonment, we should be clear, would be ruinous to the contextualists' reply to the skeptic. Without the

commitment to closure, contextualists can provide no better explanation of the lure of skepticism than rival theories offer. In order to preserve contextualism as an approach to the problem of skepticism, accordingly, contextualists must either affirm what I have denied in (2') and claim that I really *do* know that my laptop has not been stolen since I left my hotel room, or else they must deny what I have affirmed in (3') and claim that I do *not* know that my laptop is in my hotel room. But both are obviously unacceptable to the contextualist. To abandon (3') is to abandon the idea that ordinary language contexts fully support claims about the whereabouts of things on the basis of clear memories. As a semantic theory, however, contextualism may not take such an option. But to abandon (2') is simply to abandon common sense: I cannot know that a robbery has not occurred in some location where I am not present.

I am inclined to think that in order to preserve their account of epistemic standards, contextualists are going to have to count *any* new consideration that might alter our judgment of a given case as a change in epistemic context. But if that's right, then the very notion of an "ordinary" epistemic context, or even an epistemic context that ordinary speakers can be presumed to share, is otiose. Indeed, the very idea of an epistemic context becomes uselessly unstable. We cannot simply assume that some consideration that might be salient and appear relevant to one person will appear equally relevant and also be salient at a given moment to another. If simply reminding one another of such considerations changes the standards of knowledge, it follows that no sense can be made of there being any clear standards of knowledge to identify as settled by a context.

The problem, at base, derives from what is central to the contextual analysis of knowledge, namely, that it seeks to report or *describe* the standards that actually apply to S's linguistic context, rather than (as analytical epistemology has more traditionally sought to do) to *prescribe* applicable standards that may or may not be ones available or salient to S at the time of attribution. The contextualist approach presupposes that there must actually *be* a coherent set of standards applicable to S's actual linguistic context, with respect to epistemic terms, but cases such as those I have indicated seem to me to indicate that in quite ordinary situations, it is unlikely that there actually is any coherent set of standards provided by "ordinary" epistemic language-use. Actual studies of "ordinary" epistemic intuitions indicate, on the contrary, that most people actually do not have wholly consistent intuitions about what standards to apply to epistemic claims⁷. But if this is correct, there appears to be no good reason to think that a given epistemic context — at least if the very nature of a context is to be settled by actual language-use — will include a coherent set of epistemic standards by which we can consistently confirm or deny claims about what someone in that context does or does not know.

My claim in this paper is that the contextualist response to skepticism does not actually explain and provide a way to avoid the inconsistency between ordinary claims of knowledge, on the one hand, and the plausibility of skeptical attacks, on the other. The sort of example I give in (1') through (3') does not strike me as rare or few in number, and I do not see any way for contextualists to avoid the conundrum such examples seem to present, which generates inconsistency within what contextualists must count as a single ("ordinary") epistemic context.

⁷ See the survey and comment on these results in [11].

My own view about how best to respond to the skeptic is not a contextualist one, though it is a fallibilist one, for I think we really can know that we are not brains in vats or invincibly deceived, while acknowledging that what allows me to know such things does not bestow infallibility on my judgments about them. Good defenses for such a position have been offered by others, however, and would take me considerably beyond my topic in this paper. My argument here against contextualism, at any rate, if successful, shows us how *not* to be a fallibilist⁸.

References

1. Cohen, S. (1988), How to Be a Fallibilist, in Tomberlin, J. (ed.), *Philosophical Perspectives*, vol. 2, pp. 91–123.
2. Lehrer, K. (2000), *Theory of Knowledge*, Boulder, Colorado: Westview Press.
3. Cohen, S. (2000), Contextualism and Skepticism, in Sosa, E. and Villanueva, E. (eds), *Skepticism*, Boston: Blackwell.
4. Cohen, S. (2001), Contextualism Defended: Comments on Richard Feldman's 'Skeptical Problems, Contextualist Solutions', *Philosophical Studies*, vol. 103, no. 1, pp. 87–98.
5. Lewis, D. (1996), Elusive Knowledge, *Australasian Journal of Philosophy*, vol. 74, no. 4, pp. 549–567.
6. DeRose, K. (1992), Contextualism and Knowledge Attribution, *Philosophy and Phenomenological Research*, vol. 52, no. 4, pp. 913–929.
7. DeRose, K. (1995), Solving the Skeptical Problem, *The Philosophical Review*, vol. 104, no. 1, pp. 1–52.
8. DeRose, K. (2002), Assertion, Knowledge, and Context, *The Philosophical Review*, vol. 111, no. 2, pp. 167–203.
9. Cohen, S. (1999), Contextualism, Skepticism, and the Structure of Reasons, *Philosophical Perspectives*, vol. 13, pp. 58–89.
10. Cohen, S. (1986), Knowledge and Context, *The Journal of Philosophy*, vol. 83, pp. 574–583.
11. Alexander, J. and Weinberg, J. M. (2007), Analytic Epistemology and Experimental Philosophy, *Philosophy Compass*, no. 2, vol. 1, pp. 56–80.

Received: November 9, 2019

Accepted: March 19, 2020

Author's information:

Nicholas D. Smith — PhD in Philosophy, Professor; ndsmith@lclark.edu

Контекстуализм и заключение

Н. Д. Смит

Льюис и Кларк Колледж,
0615 SW Palatine Hill Rd., Портленд, Орегон, 97219, США

Для цитирования: Smith N. D. Contextualism and closure // Вестник Санкт-Петербургского университета. Философия и конфликтология. 2020. Т. 36. Вып. 2. С. 291–299.
<https://doi.org/10.21638/spbu17.2020.207>

Скептицизм сталкивает нас с парадоксом (который иногда называют «трилемма скептицизма»), в качестве примера которого рассматривается следующий: (1) я знаю, что сейчас я работаю за компьютером; (2) я знаю, что знание, что сейчас я работаю за компьютером, логически подразумевает, что я не обманываюсь или мной не манипулируют

⁸ *Contra* [1].

так, как представляет это скептическая гипотеза (это следствие называется «заключение из известного логического следствия»); (3) я не знаю или не могу знать о том, что не обманываюсь или что мной не манипулируют так, как представляет это скептическая гипотеза. Парадокс состоит в том, что эти три суждения логически несовместимы. Относительно новое направление в эпистемологии, которое называется контекстуализм, утверждает, что мы можем принять все три эти заявления в трилемме, признавая, что они не являются полностью истинными внутри одного и того же эпистемического контекста. Контекстуалисты утверждают, что мы можем знать в повседневных контекстах, но не можем знать о том, что мы не обманываемся или что нами не манипулируют в соответствии со скептическим сценарием, но последний факт истинен в ином эпистемическом контексте, чем обычное знание, которое у нас может быть. Заключение из известных логических следствий может оставаться истинным, но только до тех пор, пока рассмотренные следствия являются альтернативами, которые принадлежат одному и тому же эпистемическому контексту, что и изначальное утверждение знания. В этой статье я утверждаю, что интерпретация контекстуализмом того, как изменяются эпистемологические контексты, вместе с его принятием заключения, неправдоподобна.

Ключевые слова: эпистемология, знание, заключение, контекстуализм, фаллибилизм, логическое следствие, скептицизм.

Статья поступила в редакцию 9 ноября 2019 г.;
рекомендована в печать 19 марта 2020 г.

Контактная информация:

Смит Николас Д. — канд. филос. наук, проф.; ndsmith@lclark.edu