



UiO • University of Oslo

Contextualism, Probability and Salience

A Review of Different Forms of Epistemic Contextualism as a
Response to Philosophical Skepticism

Sander Wilhelm Gjedrem Lie

Master Essay
Philosophy
FIL4091 - 30 credits

IFIKK
Faculty of Humanities
University of Oslo

Autumn 2019

Contents:

Introduction	3
The Skeptic Threat	3
The Contextualist Solution	5
Lewisian Infallibilism	6
Fallibilism and its Forms	10
The Nature of the “Bumps”	12
Probabilistic Contextualism	13
The Lottery Problem and its Solution	14
The Need for Probabilistic Knowledge	15
Sarah Moss' Account of Probabilistic Knowledge	17
The Regress to Zero	19
Foundationism	22
Factivity and Indeterminacy	23
Knowing the Improbable	25
Salience Contextualism	26
Dust Under the Rug?	28
Truth and Appropriateness	30
Why Choose Stroud's View	32
Conclusion	33
References	34

Contextualism, Probability and Salience

Introduction

This paper concerns skepticism, and epistemic contextualism as a response to the problems posed by skepticism. I will provide an outline of the way in which contextualist notions of knowledge ascriptions attempts to solve these problems, using examples mostly from Stewart Cohen and David Lewis. I then argue that, in addition to differentiating between fallibilist and infallibilist contextualism, there is another way of categorizing forms of contextualism. While Cohen connects his contextualism to the notion of probability (where you can know p in virtue of the evidence making p probable)

The Skeptic Threat

Knowledge is something often taken for granted in our daily lives; someone asks us “Do you know what the time is?”, upon which we may glance at a watch and answer “Yes, I do. It's 12.45”. And we do seem to know quite many things, including small things, like the time, and bigger things, like the fact that Caesar conquered Gaul. We know obvious things, like basic arithmetic truths, and less obvious things, like what kind of food a certain friend enjoys. These things may be spoken of, but are generally not questioned on a day-to-day basis. There are, however, those that would dispute our supposed knowledge of these things, and more. They may dispute our knowledge of everything, such that we end up apparently knowing nothing; this is skepticism.

The skeptics have been around for quite some time, and can be traced back to ancient Greece. Here, we can differentiate between two schools of skepticism: pyrrhonian skepticism and academic skepticism. Named after Pyrrho, Pyrrhonian skepticism mainly questions our ability to apply reason to settle matters, leading to the conclusion that we do not know the extent of our knowledge (Klein 2003, p. 36). Academic skepticism, which originated from Plato's Academy (ibid.), has gained more traction in the modern era. It is mostly concerned about our inability to rule out certain scenarios that, if true, would turn our otherwise accepted beliefs false, and thus not knowledge. These scenarios, called skeptical scenarios or skeptical hypotheses, are generally hypothetical situations where “(a) its truth is inconsistent with some propositions we ordinarily take ourselves to know, and yet (b) the hypothesis is compatible with all our experience in favor of those ordinary propositions” (Hannon 2016, p. 2). For the purposes of this paper, I will consider academic skepticism the most, because of its greater prevalence nowadays, though pyrrhonian skepticism will

not be entirely irrelevant.

Exactly how, then, does the skeptic go about to undermine our knowledge? To see this, let's look to our friend René Descartes, an academic skeptic and often considered the father of modern philosophy. While sitting in his chair, writing his meditations, he realised that he often had experienced dreams so vivid they were indistinguishable from when he was awake. He could thus not be certain that he was indeed sitting in his chair, writing, since he could be sleeping. "How often does my evening slumber persuade me of such ordinary things as these: that I am here, clothed in my dressing gown, seated next to the fireplace – when in fact I am lying undressed in bed!" (Descartes 1641, p. 41). Then, it also becomes impossible for him to conclude that he actually has hands, eyes and a body, since he may only perceive them within the confines of a particularly vivid dream. This argument can be made to look more orderly like this:

1. I do not know that I am not dreaming.
2. If I do not know that I am not dreaming, I do not know that I have hands (even though I clearly perceive myself to have them).
3. Therefore, I do not know that I have hands.

An argument on this form can be made for any skeptical scenario, including us living in a simulation reminiscent of *The Matrix*, being exposed to experiments as brains in vats, or the universe being created a mere ten seconds ago. For a generalized version of the argument, we can write the following:

1. I do not know that not-SH.
2. If I do not know that not-SH, then I do not know that O.
3. Therefore, I do not know that O.

(Hannon 2016, p. 3)

Here, O stands for a piece of everyday information generally accepted as being knowledge, while SH stands for whichever skeptical hypothesis is chosen to undermine the knowledge of O. The argument goes that if knowledge is supposed to be certain, and not just glorified guesswork, we should be able to eliminate all the possible scenarios in which our beliefs turn out false. At this point, our intuitions may seem to be playing a trick on us; the premises of the argument seem

intuitively true, and the validity of the argument seems indisputable, meaning its conclusion should also be true. Yet, the conclusion seems obviously false for anyone not prone to complete skepticism. Thus, the allure of the skeptic argument as a whole is contrasted with the repulsion of its conclusion.

What makes this argument so compelling? Or more specifically, what makes premises 1 and 2 so compelling? This is closely related to what is known as the closure principle, the idea that knowledge is closed under known entailment. That is, according to the principle, if S knows that p, and that $p \rightarrow q$, then S also knows that q (although this is the formulation of the principle that I chose, there are many different ways to state it). Closure seems to be intuitively very plausible; it would seem contradictory to say that S knows she has more than five apples, but not that she has more than four. And it would be likewise contradictory to say that she knows she has hands, but not that she is not a brain in a vat on which experiments are done. After all, "S has hands" implies that she must be more than just a brain in a vat, and S, being an adept philosopher, knows that she can't have hands if she only exists as a brain in a vat. The puzzling aspect of this is that the proposition "S knows she has hands" seems true, evidenced by her seeing and feeling her own hands, while the proposition "S knows she is not a brain in a vat" seems false, since she has no evidence of this being true. However, this is incompatible with the closure principle. Thus, we are forced to draw one of three rather counterintuitive conclusions:

1. The closure principle is false.
2. S does not know that she has hands (in fact, we all know next to nothing).
3. S does indeed know that she is not a brain in a vat.

(ibid. p. 4)

The threat of skepticism, in addition to pointing out how little evidence we have for our everyday "knowledge", is that it forces us to pick one of these three evils, and when we have picked one, explain why we considered it to be an evil in the first place (ibid.).

The Contextualist Solution

This is where the contextualist project enters the arena. It is an account of knowledge ascriptions meant to save our precious knowledge from death by skepticism, and has grown to significant

popularity since its birth in the 1970s. As an account of knowledge ascriptions, it is less concerned with explaining knowledge itself, and more concerned with explaining the truth conditions of propositions on the form “S knows that p”. (That being said, explaining the truth conditions of knowledge ascriptions would probably be highly relevant to the question of what knowledge is. Thus, the question of knowledge is not entirely beside the point, and it may in fact be answered indirectly by contextualism, if the theory manages to hold water). Before we get into the specifics, the basic idea is that the truth conditions of such statements vary, depending on context. Some contexts may require absolute justification for ascriptions of knowledge to be true; these are skeptic contexts, relative to which we know hardly anything. However, relative to contexts more lenient in their requirements for justification, everyday contexts, we know all sorts of things. As an example, when someone asks you for the time, and you glance at your watch, you do, in fact, know that the time is “12.45”, because the context in which the issue is considered is not one where potential evil demons are taken into account; in order for you to know the time, it is not required that you provide evidence against any skeptical scenario; the statement “You know that it's 12.45” is true. Now, if the person asking for the time was a fan of Socrates, walking around talking to people and making them doubt their previously held beliefs, he might have brought these skeptical scenarios into light, challenging your knowledge of the current time. This would be a different context, a context in which the skeptic would be right in asserting how you don't know what time it is, because of the greater requirements for justification; “You know that it's 12.45” is false. This is the basic idea shared by contextualists, though they also differ in some ways. There are a couple of ways in which to differentiate between different kinds of contextualism, and I will in the following paragraphs present a few of these differences. After I have explained established differences between contextualists, I will present my own way of differentiating between them. This is not to say that I believe my own differentiation is more correct than what is already established, but I think there is a point to be made by exploring other dimensions of the subject.

Lewisian Infallibilism

First of all, we have the distinction between infallibilism and fallibilism, which differs in their requirement for justification in order for a knowledge ascription to be true. Let us first consider infallibilism, which is the idea that a knowledge ascription “S knows that p” is only true if S's evidence actually implies the truth of the belief p; that is, she can on the basis of her evidence eliminate all possibilities in which not-p. Infallibilism can easily be taken to imply skepticism, as it

seems to state something reminiscent of the closure principle, as S's evidence would be required to eliminate any scenario in which S does not have hands in order for her to know that she has hands; this includes scenarios like S being a brain in a vat. This apparent skepticism is witnessed by David Lewis in his paper "Elusive Knowledge", where he explains his distaste for both skepticism and fallibilism: "So we know a lot; knowledge must be infallible; yet we have fallible knowledge or none (or next to none). We are caught between the rock of fallibilism and the whirlpool of scepticism. Both are mad!" (Lewis 1996, p. 550). Lewis then makes it his mission to show that the choice between skepticism and fallibilism is a false dilemma, and to steer the ship safely between the rock and the whirlpool (ibid.). Lewis could thus be considered an infallibilist contextualist, as he would likely also consider himself. However, as we shall see, there is an infallibilist dimension to Lewisian contextualism, which we should probably expect, since his hatred for skepticism is greater than his hatred for fallibilism (ibid.), and the dilemma would be hard to circumvent entirely. In fact, it could be argued that there exists no truly infallibilistic variant of contextualism.

Because the infallibilist criterion for knowledge seem to indicate a skeptic conclusion, Lewis provides a slightly different one, with something added: "S knows that P iff S's evidence eliminates every possibility in which not-P – Psst! – except for those possibilities that we are properly ignoring" (ibid. p. 554). The idea here is that when a person knows something, there may be (or always are) certain possibilities not considered at all by the knower, they are being ignored, yet the knowledge ascription is true nonetheless. This may at first seem to imply strange conclusions, like how you can know all sorts of things as long as you ignore their alternatives. However, the word "properly" saves this idea from being quite that crazy, as Lewis sets up a list of rules for when something is properly ignored. We will consider them briefly here, though we will get back to them in more detail later:

The rule of actuality states that actuality may not be properly ignored. This ensures that nothing false is ever known; when actuality, or truth, is ignored, it is not properly ignored (p. 554-555). It should also be noted that the actual world may vary from subject to subject, as the term "actual world" is, according to Lewis, a mere indexical. When someone knows something, the relevant actuality is the actuality of the world in which that individual exists, and at that time. Thus, we can truly state that I would have known who won the football game yesterday, had I read the newspaper. In this case, the knower is not myself, but rather my counterpart in a slightly different world, who did read the newspaper (ibid.).

The rule of belief states that what is believed, or should be believed may not be properly

ignored. One cannot know something if one believes the opposite. Nor can one know something if one's evidence implies the opposite. The degree of belief in which defeats the possibility of properly ignoring may depend on the situation (p. 555-556), which further signifies Lewis' position as a contextualist.

The rule of resemblance states that if a possibility *p* may not be properly ignored, neither can any possibility that saliently resembles *p*. A function of this rule is that it defeats the Gettier cases: if you glance at a stopped clock, which just so happens to show the correct time, you do not gain knowledge of what time it is, because you have not properly ignored the possibility that it is indeed a slight different time. And because the clock has stopped, this possibility saliently resembles actuality (p. 556-557). It also solves the problem with the lottery: you cannot know that you will lose the lottery because the possibility of your ticket being a winning ticket saliently resembles actuality. "For every ticket, there is the possibility that it will win. These possibilities are saliently similar to one another: so either every one of them may be properly ignored, or else none may" (p. 557).

According to the rule of reliability, our senses, memory and testimonies are somewhat reliable, and may be presupposed to convey truth. We may (albeit defeasibly) ignore the possibility of us hallucinating, remembering falsely and people speaking falsehoods. Though this rule may be overridden by other rules, such as the rule of actuality; we cannot properly ignore possibility of hallucination if we are indeed hallucinating (p. 558). (It is quite interesting to see that Lewis admits that ignoring of possibilities may be proper while still being defeasible).

The two rules of method states that we may presuppose that a sample is representative, and that the best explanation given our evidence is true (*ibid.*).

The rule of conservatism states that we can properly ignore that which is commonly ignored by the people around us (p. 559). To avoid some possible confusion, these rules stating what you may ignore, does not implore you to ignore, but rather gives you "permission" to ignore. As such, you may know things ignored by those around you.

Finally, the rule of attention states not what may be properly ignored, but rather that when something is properly ignored, it is actually ignored. Just because a certain proposition allows you to properly ignore it given all of the previously mentioned rules, it is not properly ignored unless you are actually ignoring it (*ibid.*). Thus, Descartes did not know that he sat in his chair writing, because he was not properly ignoring the possibility that he was dreaming. This possibility may have been eligible for proper ignoring, but it wasn't actually ignored.

With this final rule, it becomes clear what Lewis means when he accuses epistemology of being “an investigation that destroys its own subject matter” (ibid. p. 550). Descartes could easily have known that he was sitting in his chair, but because he engaged in epistemology, he could not help but think he might have been dreaming, thus destroying his knowledge of sitting in a chair. In similar fashion, whenever an epistemologist considers how she cannot know that she has hands, because she may be a brain in a vat, she indeed does not know this. This would also be the case for any skeptic scenario.

So according to Lewis, we know a lot. “We know a lot” (p. 549) is actually the first sentence of “Elusive Knowledge”. We know a lot in virtue of properly ignoring skeptical possibilities all the time. Whatever you may be doing, walking, driving, cooking or watching TV, you know that you have legs, you know that the traffic light is green, you know that the soup has been simmering for four minutes, and you know that your favourite TV-show is just starting. In these situations, you won't normally consider the possibility of hallucinations, dreams or evil demons deceiving you. These possibilities are being ignored, and properly so. And as long as they remain properly ignored, they do not undermine your knowledge. But the moment a skeptic arrives and turns your attention towards these possibilities, your knowledge disappears. This is why knowledge is “elusive”; it only exists when you're not considering it. “Examine it, and straightway it vanishes” (p. 560).

When examining this notion more closely, Lewis seems to become somewhat of a closet fallibilist. Infallibilism states that in order for you to know something, all possibilities in which that something is false must be eliminated by your evidence. Lewis is a proponent of this idea, even though he does not claim that anyone has evidence eliminating the possibility of being a brain in a vat, and claims propositions incompatible with this can still be known. He is thus a fallibilist in the sense that you can know you have hands without providing evidence that you are not a brain in a vat. One may think of this as follows: when Lewis says that every possibility in which not-p must be eliminated by the knower's evidence, the term “every” is to be understood as a universal quantifier, so another way to state this would be “For all possibilities x in which not-p, S's evidence must eliminate x”. However, since the quantifier range over a certain domain, there may be possibilities not ranged over by the quantifier. This is somewhat analogous to a teacher noting that all students are present in class; it does not mean that all students in the universe are present, because the quantifier is only ranging over the relevant students, and so the statement is considered true even though there may be students in the world not present. In a similar way, a knowledge ascription of p may be true in virtue of every not-p-possibilitie being eliminated, even though some

possibilities are left; these uneliminated possibilities are not in the domain over which the universal quantifier range. But the moment a possibility is added to the domain, is incompatible with p and is left uneliminated by the evidence, p is no longer known. In admitting that there are possibilities outside the domain of the quantifier that may be left uneliminated, there appears to be a sense in which Lewis is a fallibilist, though this may be the most infallibilist contextualism that is possible without it ultimately resulting in skepticism. (It is, after all, possible to be both a contextualist and a skeptic, if you hold that no context is ever lenient enough for anyone to know anything (Hannon 2016, p. 10)).

Fallibilism and its Forms

Fallibilist accounts of contextualism are more lenient in their criteria for knowledge ascriptions to be true, and are thus less likely to be prone to skepticism. In general, fallibilists struggle less with the dilemma than Lewis does, as they prefer just to crash the ship straight into the rock in order to avoid the whirlpool. At least it is solid land to stand on. Fallibilism avoids skepticism by rejecting the closure principle, or at least some form of it. Stewart Cohen defines fallibilism as the view that: “S can know P on the basis of R even if there is some alternative to P, compatible with R” (Cohen 1999, p. 58)¹. Though rejection of the closure principle removes the problem of skepticism (it involves choosing to reject option 1 on the trilemma on page 3), it may well be a hard pill to swallow, since the principle seems highly intuitive. Now, fallibilism is not by itself a form of contextualism, as contextualist aspects only come in to the picture when you assert that what is sufficient evidence for knowledge varies with context. Fallibilism is older than epistemic contextualism, and Cohen applied a contextualist dimension to fallibilism in order to save it from its inconsistencies. While fallibilism merely claims that “S knows p ” may be true without all not- p -possibilities being eliminated by S's evidence, contextualism is what states that whatever is known by S varies with context, as there is no mention of context sensitivity in fallibilism itself.

Now, fallibilism seems to imply the seemingly strange conclusion that you can truly state the proposition “I know that p , but it is possible that not- p ” (DeRose 2016, p. 4). While it is not a blatant contradiction, it does seem to state a conjunction of two conflicting propositions. At this point, there may be differing opinions in the fallibilist camp. The question is whether or not there is

¹ This is a rejection of what Cohen calls the “entailment principle”, which Cohen states as the principle that “S knows q on the basis of reason r only if r entails q ” (Cohen 1988, p. 91). While this is somewhat different from the closure principle, they are similar, and the difference between them matters little for the exact purposes of this paper. I will therefore consider them interchangeable for all intents and purposes.

a genuine conflict between the ascription of knowledge and the admitted possibility of fault, and DeRose differentiates between what he calls intuitive fallibilism and GC-fallibilism (GC being short for “genuine conflict”). Intuitive fallibilism is the belief that a knowledge ascription can be true, while admitting some manner of fallibility with respect to that piece of knowledge (ibid., p. 1). This may be somewhat vague, as DeRose admits, but it is possible to grasp the basic idea; according to intuitive fallibilism, everyday ascriptions of knowledge are true, even though the basis of that knowledge is fallible, and thus there is not necessarily a conflict between the knowledge that *p* and the possibility that not-*p*. On the other hand, GC-fallibilists admit that there is a genuine conflict between stating the statements “I know that *p*” and “it is possible that not-*p*” when stated by the same individual within the same context. DeRose understands this as relating to epistemic risk or chances of error, such that the infallibilist holds that *S* only can know *p* if there is no chance of error in respect to *p* from *S*'s point of view, while the fallibilist denies just that (DeRose 2016, p. 4). However, DeRose points out that on this interpretation, some fallibilists (himself included) may be incorrectly categorized as infallibilists. On DeRose's view, *S* cannot truly state that not-*p* is possible if she knows that *p*, and thus the statement “I know that *p*, but it is possible that not-*p*” expresses a genuine inconsistency (ibid.). This may make the term “GC-fallibilist” seem somewhat oxymoronic, as the “GC” implies that you cannot know *p* when there is a chance that not-*p* while “fallibilist” implies that you can. DeRose's solution to this problem is that one can consider “*S* knows that *p*” inconsistent with “it is possible from *S*'s perspective that not-*p*” when the standards are held constant, which still allows for *S* to know *p* according to some epistemic standard while it is possible that not-*p* according to a different epistemic standard. The meaning of epistemic modal statements and knowledge ascriptions thus “sway together” (ibid.). Then the statement “*S* knows *p* according to epistemic standard *x*” is inconsistent with “It is possible from *S*'s perspective that not-*p* according to epistemic standard *x*”, but is not necessarily inconsistent with “It is possible from *S*'s perspective that not-*p* according to epistemic standard *y*”.

To illustrate his point about epistemic risk, DeRose proposes an analogy to the notion of bumps on a surface. Often, we call surfaces “flat” even though they have bumps which would make them non-flat (ibid. p. 6). For example, a road may be flat, but there are still obvious bumps on it when you examine it. A glass window, which is much flatter than a road still has “micro-bumps” (ibid.) that are invisible to the naked eye. Still, we deem these surfaces flat despite their bumps. Similarly, an ascription of knowledge can be considered true even where there exists “micro-risks” of error. And

just like with the flatness of surfaces (If a window was as flat as a road, we would not deem it flat), the truth value of knowledge ascriptions depends on the context; in some contexts, otherwise noticeable risks do not destroy the knowledge, while other contexts may tolerate only “micro-risks” of error. In the most skeptical contexts, no risks are tolerated, which means next to no knowledge ascriptions are true.

The Nature of the “Bumps”

Thus far, it may seem all well and good; it seems intuitively plausible that the ramblings of a skeptic are in many cases misplaced and irrelevant. If not, getting all questions correct on a school test would be easy when you could answer all questions correctly with “I don't know because we may be living in a sceptical scenario”. When most people speak of what they know, the mentioning of skeptical scenarios will more often than not be beside the point and cause people to roll their eyes (For example, think of Lewis' example of sceptical scenarios mentioned in court (Lewis 1996, p. 560)). In these cases, intuition may tell us that the skeptic has a point, but that it is irrelevant in this particular situation (you asked for the time, not what can fool you into determining the time wrongly). Skepticism then seems to be in some sense true, but only in the appropriate context. The conclusion is that context is somehow relevant to the truth of knowledge ascriptions. This appeals to our intuitions both favouring skepticism and those favouring everyday knowledge ascriptions, making the prevalence of contextualism unsurprising.

A common feature of all forms of contextualism is the idea of some propositions, while not refutable, need not be refutable in all contexts because they would be irrelevant in some contexts; indeed, Stewart Cohen does not use the term “contextualism” when forming his theory in “How to be a Fallibilist”, rather, he uses the term “theory of relevant alternatives” (Cohen 1988, p. 94). Here, however, I would like to pose a question to the contextualist; not an objection per se, but a request for clarification. Because regardless of which form of contextualism we speak of, we may still wonder what kind of parameter determines the standard of a given context. How does one context differ from another? Or more specifically: Exactly what is required by a strict (or skeptical) context but not by a more lenient context? In order for contextualism to work, it needs to account for the nature of one context as opposed to another; if a particular “bump” in everyday knowledge is considered “small” enough not to render the knowledge ascription false in a given context, we must have a way of determining the “size” of the bumps. Thus, we must ask the question of what unit of measurement makes up that size. Now, all forms of contextualism agree that a skeptic uttering

something like “But what if we are brains in vats?” changes the epistemic landscape in a certain way, such as to make the truth conditions of knowledge ascriptions harder to satisfy. But what has changed? When an irrelevant possibility suddenly becomes relevant, what made it irrelevant before, and what made it relevant now?

In my view, this is where epistemic contextualists disagree with each other in a way the fallibilist-infallibilist distinction does not capture, or even tries to capture. The most important way in which Cohen differs from Lewis is not that Cohen is a fallibilist and Lewis is an infallibilist, at least not for the purposes of this paper. In order to lay the groundwork for my critique of epistemic contextualism, I must first account for Cohen's and Lewis' answers to the question of what determines the size of “bumps” in our knowledge. They both have their own problems, but Cohen's problems are ultimately greater than Lewis' problems. Thus, Lewis' account remains more plausible in the end, though necessarily plausible enough to account for how knowledge ascriptions can be true in the face of sceptical scenarios.

In the following sections, I will propose a new way of differentiating between various forms of contextualism; instead of considering the distinction between fallibilism and infallibilism, I will focus on the question of what makes a certain context different from another, and what constitutes the “strictness” of a given context. My critique of these forms of contextualism will then be based on what I discover about their accounts of this particular issue. First, I will argue that Cohen's reference to probability requires us to be able to know probabilistic contents, and that this ultimately causes his project to fail, mostly because of the regress problem, discussed in the section “the regress to zero”. I will then argue that Lewis' salience account is not subject to the same problems, and is thus more plausible (though still subject to more general issues with contextualism).

Probabilistic Contextualism

The first form of contextualism I would like to consider is what I would like to call probabilistic contextualism, in which knowledge ascriptions of less certain propositions can still be true in virtue of the truth of the proposition being probable. The way to understand strictness of contexts would then be that stricter contexts require higher probability for the relevant proposition p in order for “ S knows that p ” to be considered a true knowledge ascription. In more specific terms, every context's standard would have a threshold number (T) set between 0 and 1, and in order for “ S knows p ” to be true, S 's evidence must render $P(p)$ (the probability of p) equal to or greater than the threshold. In a skeptic context of maximum strictness this threshold number would be no less than 1, a standard

expecting absolute certainty from S's perspective. According to any standard for evaluating knowledge ascriptions, "S knows p" is true just in case $P(p) \geq T$, $P(p)$ being the probability of p from S's evidence, and T being the threshold set by the standard.

One who seems to belong to this school of contextualism is the already mentioned Stewart Cohen, judging from the beginning of "How to be a Fallibilist", where he states that "a fallibilist theory allows that S can know q on the basis of r where r only makes q probable" (p. 91). He speaks here of fallibilism rather than contextualism, but introduces his theory of relevant alternatives as it becomes clear that probability in itself can not be the sole determining factor in whether or not a somewhat uncertain proposition can be known, as this brings problems. Keith DeRose also has formulations that hint towards a notion of probabilistic contextualism, since he talks of "a notion like that of epistemic risk, or chances or possibilities of error to draw the distinction" (DeRose 2016, p. 4). DeRose's references to probabilities are less explicit, so I will mostly be considering Cohen here.

The Lottery Problem and its Solution

At this point, there is an elephant in the room, a common objection to the claim that we can know things in virtue of them being probable; this is the lottery problem. While the objection does show that probability cannot be the sole decider of knowledge, I believe Cohen successfully defends against it for the purposes of contextualism. I will address the objection here to show my reader that probabilistic contextualism is not defeated easily by the obvious objection. It goes as follows:

Suppose you have bought a ticket for a lottery; given the unlikelyhood of your ticket being a winning ticket, may you invoke contextualism in order to truly state that you know you won't win the lottery? This does not seem right, as there are a finite number of tickets and all tickets have an equal probability of winning, unless the lottery is rigged. Therefore, it does not seem allowed to disregard the chances of any single ticket, as one ticket is bound to win, and it could be any of them. This is also the case for arbitrarily large numbers of tickets, meaning we can make the probability of winning arbitrarily small without permitting S to know q (q being the proposition that S loses the lottery). However, according to Cohen, it is possible to know you will lose the lottery if it has come to your attention that the lottery is rigged such that you will lose (Cohen 1988, p. 92). This is where context becomes relevant to the ascription of knowledge. As mentioned, it is possible to make the probability of loss at a fair lottery arbitrarily close to 1, without making it possible to know you will lose (ibid.). On the other hand, if you play a small lottery of 10 participants, and have it on good

authority that the lottery is rigged against you with 99% certainty, it is possible to truly state that you know you will lose. However, the actual probability of loss is a mere 0,999. This may seem a high probability, but it is still much lower than the probability of loss at a lottery with billions of tickets. Yet, even though the chance of loss is lower, it is only possible to know you will lose in the case of the rigged lottery, not in the fair lottery.

It then becomes clear that probability by itself can not be the sole decider of what can count as knowing and what can not, though it may still be one of the deciders. According to the theory of relevant alternatives, as Cohen calls it, some alternatives in some contexts can be, and indeed are, disregarded as irrelevant. In the case of a fair lottery, one can never disregard the alternative of any particular ticket winning, as all tickets are equally relevant. However, in the case of the rigged lottery, all tickets are not equally relevant, and even though there is a chance you may be mistaken about the lottery being rigged, or the rigging may fail, these alternatives can in certain contexts be disregarded as irrelevant. In this case, these alternatives are unlikely enough to be justifiably disregarded in this context. In the context of the fair lottery, the very unlikely alternative of you winning is still not unlikely enough to be disregarded. This is how Cohen escapes the lottery problem. In his own words, "How probable an alternative must be in order to be relevant will depend on the context in which the knowledge attribution is made" (Cohen 1988, p. 96). Thus, the context of a fair lottery is never such that the standard allows for S to know that she loses.

The Need for Probabilistic Knowledge

However, while Cohen may have avoided the lottery problem, I must here present my first objection to Cohen's account (the lottery problem is after all not of my own making), the need for probabilistic knowledge. Consider this: If it is possible for "S knows p" to be true in virtue p being probable enough from S's evidence, as the evidence puts $P(p)$ at or above the threshold number T set by the standard of the context, S will also need to have knowledge of $P(p)$ as higher than T . In other words, in order for "S knows p" to be true, "S knows $P(p) \geq T$ " must also be true. If this is true, we would need to account not only for how we can know that we have hands, but also for how we can know that we probably have hands. And the same would go for every other knowledge ascription.

"Slow down!", you may say, because Cohen speaks of S knowing p, not S having any knowledge of $P(p)$. Wouldn't it be enough, as Cohen says, that p is probable enough according to

the standard set by the context, according to S's evidence? Possibly, but what does that really mean? Hear me out: First of all, it is clear that if "S knows p" is true according to probabilistic contextualism, S can not regard p as too unlikely to be concluded. S must therefore believe that $P(p) \geq T$ in order to know p. This should be uncontroversial, and fully consistent with Cohen's own view. Secondly, if S is to know p according to probabilistic contextualism, $P(p) \geq T$ must be justified by S's evidence, as noted by Cohen:

If S were to have no evidence at all regarding the behavior of zoo keepers (or even people in general) which would count against this alternative-if as far as S's evidence were concerned, it would be as likely as not that he sees a disguised mule-then surely S would fail to know that he sees a zebra (Cohen 1988, p. 102)

Thirdly, I claim that, if probabilistic contextualism is to work, S cannot know p without $P(p) \geq T$ being true. Consider the following example, somewhat akin to the Gettier cases: If S believes she has won the lottery, based on a generally reliable source that the lottery is rigged in her favour, this belief might constitute knowledge according to epistemic contextualism. However, if the lottery is indeed not rigged at all, S does not know that she has won the lottery, even if her ticket happens to be a winning ticket. Being right about p is not enough; S must also be right about probably being right that p; if not, S's alleged knowledge that she wins the lottery turns out as epistemic luck, which is not knowledge.

It turns out then, that $P(p) \geq T$ must be a justified true belief from the part of S if "S knows p" is to be true, and I thus conclude that it must also be known. Now, as some critics of my conclusion may point out, the Gettier cases has shown that it is possible to have a justified true belief without having knowledge. So, perhaps S need not know $P(p) \geq T$, but merely have a justified true belief. Though I can not refute this notion, it does strike me as rather odd, and possibly a bit dangerous. Because the Gettier cases show how justified true belief can fail to constitute knowledge because of epistemic luck, so if S justifiably and truly believe that $P(p) \geq T$ without knowing it, it will turn out as luck. If S then concludes p on because of $P(p) \geq T$, the belief that p will also become epistemic luck, or at least be dangerously close. The cases where $P(p) \geq T$ can be justified true belief without knowledge are too few and too specific to be of any use to Cohen's account, and if they prove to be of use, epistemic luck still seems to undermine any knowledge produced by $P(p) \geq T$. I conclude then, that probabilistic contextualism can not work unless knowledge of probabilistic contents is possible and indeed real; specifically, probabilistic contextualism requires

that “S knows p” implies “S knows $P(p) \geq T$ ”. It thus behooves Cohen to have an account of how probabilistic contents can be known.

Sarah Moss' Account of Probabilistic Knowledge

Now, I will digress for a bit, as this section does not directly concern Stewart Cohen or his contextualism. I do however, think that because of my conclusion in the previous paragraph, this ultimately becomes relevant to Cohen's particular form of contextualism. What I would like to discuss, is Sarah Moss' account of probabilistic knowledge. These paragraphs up until the start of the section “the regress to zero” will therefore not concern Cohen directly, and will therefore not contain any objections to his ideas. On the contrary, I think that if anything can save Cohen from by previous objection, it would be an account like the one presented by Sarah Moss. Therefore, I will now use some time outlining the basics of her account, and point out how this helps Cohen.

Because as I concluded previously, Cohen and other proponents of probabilistic contextualism would benefit from an account of how one can know probabilistic contents. As it happens, such an account does indeed exist, provided by the philosopher Sarah Moss. In her book, aptly named *Probabilistic Knowledge*, she argues that we can believe, assert and indeed know probabilistic contents, in addition to full beliefs (Moss 2018, ch. 1, p. 2). Probabilistic contents can be expressed by sentences like “it will probably rain tomorrow”, “I might have the flu” or “there is a 0.6 probability that Liverpool will win the game tonight”. Thus, they generally have a certain resemblance to modal statements. Moss writes:

For example, say you believe your friend Smith smokes, while you have .6 credence that Jones smokes and .3 credence that Brown smokes. Just as your full belief that Smith smokes can be knowledge, your .6 credence that Jones smokes can be knowledge, and so can your .3 credence that Brown smokes (ibid.)

If Moss is right, “S knows p” can be true also if p states probabilistic contents. The previously mentioned $P(p) \geq T$ is an example of probabilistic contents, meaning “S knows $P(p) \geq T$ ” can be true according to Moss. This renders probabilistic contextualism plausible in the face of my objection from the need for probabilistic knowledge. As an example, suppose S has a 0.98 credence that p is true; according to Moss, this credence can be an instance of knowledge. What this matters to Cohen is that in a context where $T = 0.97$, it is possible that “S knows p” is true in virtue of “S knows $P(p) \geq T$ ” is also true. If a skeptic comes along and raises the standard to $T = 0.99$, “S knows

p” is no longer true. This is in complete accordance with what Cohen is trying to achieve.

How then, can we come to have knowledge of such probabilistic contents? According to Moss, it is through all the same means as how we attain “normal” knowledge, or knowledge without reference to any probability (i.e. “Jones smokes”). That is, through “testimony, perception, inference, memory, and a priori reflection” (ch. 3, p. 3). A meteorologist's testimony about how the weather will probably be, may lead to a high credence that the weather will be like that; remembering that Jones swims every tuesday may lead to a high credence that he will swim this tuesday. These things may form probabilistic content beliefs, and in many cases may in Moss' view constitute knowledge. Note that remembering Jones swimming every tuesday, and thus believing he will probably swim this tuesday, does mean you know he will swim this tuesday, merely that he probably will. In order for probabilistic contextualism to work, knowledge that probably p must in some cases, in certain contexts, imply knowledge that p.

There are a few things to point out when it comes to Moss' idea of probabilistic knowledge. First of all, probabilistic contents are different from propositional contents. One might be tempted to think of a 0.5 credence that Jones smokes as a full belief in the proposition that there is a 0.5 chance that Jones smokes, thus making credences reducible to full beliefs. This is not the case according to Moss. While propositions may be considered sets of possible worlds (their members being those worlds in which the proposition is true), “The content of a probabilistic belief is a set of probability spaces rather than a set of worlds, and so it is not true at a world simply in virtue of containing that world” (ch. 6, p. 4). Thus probabilistic contents are not propositions and beliefs in them cannot be reduced to full beliefs. In Moss' view, it is the other way around; full beliefs are reducible to credences. This assertion may tempt us to define a full belief as a credence of 1, but that would render full belief the same as absolute certainty, which in turn would mean that we have fewer full beliefs than we actually do. Instead, Moss' account of full beliefs as probabilistic contents are connected to the notions of strict and loose speech. When we utter the phrase “it is 3 o'clock”, we hardly mean to claim accuracy to the nanosecond. Normally, we speak loosely, and thus the statement can be true also at 3.01 (Moss 2018, ch. 3, p. 22-23). In the same way, “Jones smokes” can be uttered if one merely has a very high credence that Jones smokes. “Asserting that Jones smokes leads to our accepting for purposes of conversation that it is certain that Jones smokes, which is why it sounds bad to follow up by saying that Jones might not smoke after all” (ibid., p. 23). These statements by Moss' does seem to have a resemblance to Cohen's account of

contextualism, though Moss at this point speaks more of the content of utterances than the nature of knowledge and knowledge ascriptions.

Another thing to point out is that Moss herself uses contextualist reasoning to defend her own account against possible objections from skeptics. A potential problem here is that if one wants to use Moss' account to argue in favour of Cohen's account of contextualism, we could end up in a situation in which the accounts of probabilistic contextualism and probabilistic knowledge rest on each other's shoulders. Thus, a possible objection to Moss' account is that in defending her account against skeptical arguments, she turns her framework circular. However, it is not clear that Moss argues that "S knows p" can be true in virtue of S's evidence making p probable enough in a given context, and thus she may not be a contextualist of the probabilistic school. She merely accepts contextualism about full beliefs as an effective way of combating skepticism, and then shows how the same contextualist arguments apply to probabilistic contents (ch. 7, p. 41). For this reason, and because it would be slightly off topic for the purposes of this paper, I will give Moss the benefit of the doubt, and not accuse her of circular reasoning in this case. For now, let it just be said that Moss seems to be an epistemic contextualist herself; she may not be a probabilistic contextualist, and hopefully she is not, as this could render her account circular.

Since this is not a paper mainly concerning Sarah Moss' probabilistic knowledge, I will leave her be for now. The important thing to note is that Moss shows that sentences on the form "S knows $P(p) \geq T$ " can plausibly be true. And as I concluded in the previous section, Cohen needs this to be the case. But even though Cohen may be momentarily saved by Moss, I still have a few objections to his account, which make it inferior to Lewis' brand of contextualism. I will present these objections in the following sections. They concern the regress problem caused by my earlier objection, the factivity of probabilistic contents, and the issue of whether or not the things we take ourselves to know are indeed probable, as sometimes it seems they are not.

The Regress to Zero

The objection I have called "Regress to Zero"-objection is a result of my conclusion that probabilistic contextualism implies that knowledge of probabilistic contents is required for knowledge of anything. I will now argue that this ultimately makes knowledge impossible, due to an infinite sequence of products approaching a total probability of zero. I will also consider certain possible solutions to this problem, and argue that they are ultimately unsuccessful.

Now, if “S knows $P(p) \geq T$ ” must be true if “S knows p” is to be true, as I have argued is the case for probabilistic contextualism, this gives rise to a particular problem. If we abbreviate the statement “ $P(p) \geq T$ ” to simply “a”, we end up with the statement “S knows a” being a prerequisite for “S knows p”. But this would also beg the question of how S can be said to know a. According to probabilistic contextualism, S knows a in virtue of S's evidence making a probable enough in the given context, that is the probability of a is equal to or greater than the threshold set by the standard of the context, or $P(a) \geq T$. Then, in the same way as before, “S knows $P(a) \geq T$ ” must be true in order for “S knows a” to be true. The same requirements for “S knows p” also applies to “S knows $P(p) \geq T$ ”. So you cannot even know $P(p) \geq T$ without also knowing $P(P(p) \geq T) \geq T$, and as long as one holds on to probabilistic contextualism, this regress is bound to continue.

This problem is looking very similar to the epistemic regress problem, which challenges the idea that human knowledge has a foundation. “For a belief to be justified it is not enough for it to be accepted [...] there must also be good reason for accepting it. Furthermore, for an inferential belief to be justified the beliefs that support it must be justified themselves” (BonJour 1978, p. 107 (quoted from Quinton)). The question becomes: does this regress stop somewhere, and we reach belief that do not owe its justification to any other (ibid., p. 1080), or does it continue indefinitely, never reaching a foundation? The epistemic regress problem may not be entirely the same as the regress problem for probabilistic contextualism, but they are at the very least similar. In fact, as we shall see, the problem with probabilistic contextualism is arguably worse. But at least we can try to use the similarity of the two issues to search among solutions to one for a solution to the other, which I will do in the next few paragraphs.

Out of all the attempted solutions to the epistemic regress problem, infinitist solutions are the ones I find the most promising, so I will start with those. Unfortunately, this is where we come across a particularly cheeky aspect of the probabilistic version of the regress problem not found in the original. The problem is that S, in order to know p, must justify p by $P(p) \geq T$, which again must be justified by $P(P(p) \geq T) \geq T$, and so on, and all these statements must also be known by S. If any one of these statements fails to be known by S, they all fail in this regard. In other words, if we call “S knows $P(p) \geq T$ ” the predecessor of “S knows p”, any statement “S knows p” hinges upon its predecessor “S knows $P(p) \geq T$ ”, and indeed all predecessors beyond. In a similar way, from S's perspective, p hinges upon $P(p) \geq T$. And since this is probabilistic contextualism, all these statements are true in virtue of S's evidence making them probable enough for the context. Their probabilities are all less than one, since there may always be possibilities deemed irrelevant in the

context, like a zebra in a zoo in reality being a disguised mule. Suppose then, that S has concluded p on the basis of evidence. There might be a chance of error, but this is small enough not to be undermining, that is, S (consciously or otherwise) concludes $P(p) \geq T$. Still there may be a chance of error, but this is also too small to be an issue; $P(P(p) \geq T) \geq T$. This continues, and every time, we because the statements hinges upon their predecessor, we must multiply their probabilities together, and the result we get from this will be the probability that S's belief that p constitutes knowledge: $P(S \text{ knows } p) = P(p) * P(P(p) \geq T) * P(P(P(p) \geq T) \geq T) * \dots$ and so on. As for the result of this product, it becomes something like the limit of n^x , n being a number close to but not quite 1, as x approaches infinity (though the exact value of n may in this case vary for every iteration of x). As you may be aware, the product of infinitely many numbers less than 1, comes out as zero. This means that as long as “S knows p” hinges upon “S knows $P(p) \geq T$ ”, knowledge is impossible, since it turns out the likelihood of S's beliefs being knowledge as zero. S may still have true beliefs, but these would then be results of epistemic luck. This applies even when all values of n are arbitrarily close to 1, as long as they are not exactly 1. Thus, an infinite regress of this sort undermines S's knowledge completely, making all statements “S knows p” false, perhaps save for some necessary truths, in which case we are back to where we started: the skeptic position.

One could try to solve this problem by adopting a stance where the regress is avoided by way of provisional justification, admitting that the regress continues, but accepting some statement in the structure tentatively, without reference to a predecessor. Such a solution to the epistemic regress has been presented by Peter Klein in “Human Knowledge and the Infinite Regress of Reasons”. Klein is an infinitist according to whom all justification is provisional (Klein 1999, p. 178). S may believe p on the basis of reason r, r thusly serving as provisional justification for p. This is possible even without provisional justification for r, meaning we need not continue the regress beyond r. Now, should the need arise, we can find provisional justification for r as well, but this is not needed in order to provisionally justify p. The idea is that at any point in the structure of regress of reasons, we can find provisional justification for our beliefs, a sort of provisional foundation. The regress does continue beyond it, and we can provide justification for the provisional foundation whenever we need or want to.

I would not characterize our epistemic predicament as one in which there are not nets. For there might be a net whenever we need one. Rather, I would characterize it as one in which it is possible, as Lewis Carroll would say, that there are nets all the way down (ibid.).

This may work as a solution to the classic epistemic regress problem, but the regress of probabilities is a different story. For by accepting the idea of provisional justification, we must also concede that there is indeed an infinite regress; we just do not have to consider all of it at all times (actually, we never have to consider all of it). This means that if we base our justification on probabilities, the regress still brings us down to a probability of zero, as our provisional justification does not stop the regress. As such, we need the regress to actually stop, and find a foundation that is more than provisiona.; this brings us to foundationism.

Foundationism

Foundationism is the idea that there is not an infinite regress of reasons and justification. At some point, we reach a foundation that is not justified by further beliefs. One way to consider this in relation to probabilistic contextualism is that we reach a point where one of our beliefs has a probability of 1, making error impossible. For example, S can conclude from the evidence that $P(P(p) \geq T) = 1$. However, within the confines of Cohen's contextualism, this S still knows this in virtue of evidence making it probable, meaning there is still a chance that $P(P(p) \geq T) = 1$ is false; S's evidence merely make it probable. Because of this, the regress actually does not stop, since you would still have to justify $P(P(p) \geq T) = 1$, in virtue of the evidence making it probable. The only way to avoid the regress to zero in this way would be if at some point, all probabilities beyond that point were equal to 1. This seems dubious, and probably contrary to what Cohen actually thinks; it would require certainty from an account that means to explain knowledge without the need for certainty, rendering the project self-defeating.

Another way to consider foundationism is to say at some point in the regress, we will reach a belief that does not owe its justification to another. In this case, it means that the belief, i.e. $P(p) \geq T$ is not justified by S's evidence making it probable, as this would continue the regress. There may be some feature of the claim, F^2 , which justifies $P(p) \geq T$. This is called a meta-justification (Klein 1999, p. 169). This view is still problematic; first of all, Klein points out that "Either the meta-justification provides a reason for thinking the base proposition is true (and hence, the regress does not end) or it does not (hence, accepting the base proposition is arbitrary)" (p. 170). It thus seems very difficult to properly found any claim without further reason. Another problem with this construal, if used as a defence of Cohen's contextualism, is that it actually constitutes an abandonment of Cohen. Cohen answers the skeptic by stating that S can know p as long as the

2 Originally, the letter P is used. To minimize chances of confusion for the reader, I switched to F because of the otherwise prevalent use of the letter P in this paper.

evidence makes p probable. Now, S must know p in virtue of something else, namely that p has the property F . The account is then rendered insufficient, as other ways of making knowledge ascriptions true would be needed. In that sense, the foundationalist approach to the regress problem makes the problem even bigger, as it demands an entirely new account of how S can know p .

Factivity and Indeterminacy

Let us leave the regress problem for this time, and move on to my next objection to probabilistic contextualism, which concerns the factivity of probabilistic contents. As “ S knows p ” is generally accepted to imply “ p ”, there must be a way in which p can be true. And since Cohen's account requires certain probabilistic contents to be true, these must have a truth value; there must be a sense in which statements concerning $P(p)$ are true or false. On the face of it, this also seems to be definitely the case. For example, I may believe that when rolling a particular six-sided die, the probability of rolling a six is $1/6$, which would be the statement $P(x = 6) = 1/6$. And as long as we are talking about fair dice, this statement is true. On the other hand, if the die is rigged to produce either less or more sixes, the statement is false. So in this particular case, there seems to be no problem in considering probabilistic contents as factive. There are also other cases like this: the chance of a newborn baby being female is (very close to) 0.5 . The chance of winning a fair lottery with 100 tickets and one winning ticket, by buying one ticket, is 0.01 . These cases are all quite easy to handle because they can all be examined statistically, by performing 10 000 die rolls, or counting the number of females among 10 000 babies. The fair lottery can even be considered analytically, where we can conclude that out of a hundred equally likely outcomes, one outcome results in a win, giving the result 0.01 . These two methods, analytically or statistically, is how we determine probabilities. They also seem to, in a certain sense, describe what it means for the respective probabilistic contents to be true. The chance of a baby being female is factively 0.5 because half of all babies are female, the chance of winning this lottery is factively 0.01 because there is one possible success among 99 possible failures, all equally likely.

However, not all cases are as straightforward as this. For what is the probability that Liverpool wins the game? Is it likely that Jones smokes? And what is the chance of the external worlds existing? No matter what one may believe about these statements, it seems that what governs the probabilities of these cases are far more complex than those mentioned in the previous paragraph. According to Sarah Moss, probabilistic contents are indeed factive, and the belief that Jones probably smokes “is true just in case Jones probably smokes” (Moss 2018, ch. 6, p. 4). This

seems trivial, being an instance of something like the conditional $p \rightarrow p$, and so it does not explain how the statement can be factive. She provides a somewhat better explanation in the footnotes of the same chapter: “To be more precise, one could say that a probabilistic content is true just in case it contains the probability space that assigns probabilities to all and only those propositions that are objectively possible, and which assigns each proposition its objective chance” (p. 14). This may better explain how true probabilistic contents relate to objective chance, but does not explain the actual nature of objective chance. When it comes to the objective chance that Jones smokes, it becomes more difficult to construe. In a deterministic world, it may be easy, because all outcomes would, given the exact state of the world and the laws of the world, have objective chance either 0 or 1; everything is bound to occur or bound not to. Moss is clearly not such a determinist, since she believes in chances other than 0 and 1. It seems hard to explain how probabilistic contents not analyzable statistically or analytically have an objective chance. One could perhaps say that “Jones probably smokes” is true statistically, by looking at different characteristics such as age, gender, nationality, social standing, the fact that you saw her smoke last week etc., and make a statistic out of this, and calculate the likelihood of Jones smoking. However, the more variables you take into account, the more accurate your result will be. Thus, if you take ALL variables into account, the result will be completely accurate and reflect the objective chance of Jones smoking. But since no two people are completely the same, Jones will be the only individual in the statistic, making the chance of Jones smoking either 0 or 1, depending on whether or not Jones really smokes. Not only does this give an unwarranted result, it also requires you to know whether or not Jones smokes in order to know whether or not she probably smokes, which is even more unwarranted. The problem with the factivity of probabilistic contents is not so much the problem of relating probabilistic belief to objective chance, but more to explain the nature of objective chance itself. Whether or not the chance of Jones smoking is 0.6 or 0.7 does not seem to make an actual difference. What makes the difference is whether or not Jones indeed smokes, and what grounds Smith would have to believe that she smokes. The only difference probability makes in this case is to the state of Smith's mind, meaning Smith's belief is not really true or false in the traditional sense. Some probabilistic contents thus seems indeterminate, and thus not truly factive. This poses further problems for probabilistic contextualism, who require probabilistic contents to possibly be known; if you cannot know that Jones probably smokes, neither can you know that she does smoke.

Knowing the Improbable

Before we leave probabilistic contextualism, I would like to briefly discuss one last issue, namely the matter of to what extent the things we “know” are actually likely. In his paper “Very Improbable Knowing”, Timothy Williamson argues that “S knows that p” can be true even when the probability of “S know that p” being true, from the perspective of S's evidence, is arbitrarily close to zero (Williamson 2013, p. 2). I will not have the time to go into depth of Williamson's thesis, but let it be said that it is not obvious that knowing is necessarily connected to high probabilities. Now Cohen speaks of “S knows p” being true in virtue of p being likely, not “S knows p” being likely, meaning they are not the entirely same case, but the point remains that “S knows p” can be true even though S's evidence deems it very unlikely, seemingly contradicting Cohen's thesis, or at the very least challenging it.

There is also an argument out there called the “simulation argument”, which concludes that we probably live in a simulation (à la the Matrix). Nick Bostrom argues that at least one of the following statements must be true: (1) Humans will likely go extinct before we develop the technology to run detailed simulations. (2) Civilizations capable of running such simulations are not likely to. (3) We are very likely to live in a simulation (Bostrom 2001, p. 1). Again, I do not have the time to go into the details of this argument, but I mention it to make it clear that there are arguments that skeptic scenarios are not only impossible to refute, but actually very likely to be the case. I do not say this to condone the simulation argument³, but to make the point that just because something is a skeptical scenario does not make it necessary unlikely, as there are arguments to the contrary. In fact, similarly to what was discussed previously, skeptical possibilities are very difficult to dismiss on the basis of their improbability, as that seems incalculable. How likely is it that you are being deceived by a demon? How would you even begin to determine that probability? Cohen does have an answer to this. The case of the deceptive demon is what he calls a radical skeptical hypothesis, which is “designed to neutralize any evidence that could be adduced against them” (Cohen 1988, p. 111). Cohen admits that there are no evidence against such a possibility, but states that it can still be rationally believed because it is “intrinsically rational” (p. 112). He also writes: “While the skeptic may have an argument that we possess no evidence against radical skeptical hypotheses like the demon hypotheses, he has no argument that it is not rational (in the way I have indicated) to deny them” (p. 113). He may be right that there is no such argument from the skeptic, but the same argument could be made the other way around; just as there is no argument

³ I actually think that the simulation argument is fallacious, but that would be a story for another time. However, it serves to show that the things we take for granted are not obviously probable.

that denying a radical skeptical hypothesis is irrational, there is no argument that denying a non-skeptical hypothesis. For example, there is no argument that denying the existence of the external world is irrational. The idea of intrinsically rational beliefs thus becomes somewhat subjective, and if non-skeptic assumptions may lead to true knowledge ascriptions, so may skeptic ones, in certain cases.

All in all, I think the problems with Cohen's theory of relevant alternatives, or probabilistic contextualism, has too profound problems that are not easily solved. I will now discuss an alternative way of construing the contextualist project, which is David Lewis' form of contextualism. I will argue that it is superior to Cohen's theory, because it avoids many of the problems Cohen encounters without having any major problems Cohen does not have.

Salience Contextualism

Cohen thinks some possibilities are irrelevant in virtue of them being too improbable relative to the standard provided by the context, and thus these possibilities can be ignored. Lewis instead uses the “rules” mentioned earlier in this paper, the rules of actuality, belief, resemblance, reliability, method, conservatism and attention. As long as we follow these rules for what we can and cannot ignore, we can safely ignore many propositions coming as a result of skeptical hypotheses. Lewis' account is one of salience⁴, and I will call it Salience Contextualism. What makes a proposition relevant to Lewis, is whether or not it is salient. And what determines the salience of a proposition is not its likelihood, but its membership in the domain of salient propositions. When we say “I know p because I have eliminated all not-p possibilities”, we are not talking about ALL possibilities, in the same sense that a teacher in class who notes that “everyone is here” is not talking about EVERYONE. Thus, when the skeptic says “but you haven't eliminated the possibility of a deceptive demon, so you haven't eliminated all not-p possibilities” it is the same as if one of the teacher's students would say “but Queen Elizabeth II is not here, so not everyone is here”. Skeptical scenarios are simply not in the universe of discourse, and thus not captured even by the universal quantifier. This also means that, according to Lewis, relevance of possibilities are dichotomous. In Cohen's account, a proposition can be irrelevant, but closer to being relevant than a certain other irrelevant proposition, if it is more probable. This is not the case for Lewis; a proposition is salient (and thus relevant) if and only if it is in the universe of discourse. It is also clear that Lewis is not a

⁴ Arguably, Cohen's theory is also one of salience, but it would be salience in virtue of something like justification and/or probability. This is merely a question of terminology, and I chose this terminology.

probabilistic contextualist since he dismisses this idea entirely on the basis of the lottery problem. In fact, he dismisses the idea of justification as a key component of knowledge, claiming it is neither sufficient nor necessary (Lewis 1996, p. 550-551). This brings me to the first reason why I deem Lewis' contextualism more successful than Cohen's: Hilary Kornblith makes a distinction between two kinds of skeptic; the High Standards Skeptic and the Full-Blooded skeptic:

The reason I don't know anything about the external world, according to the skeptic, is not that I have a small degree of justification for my beliefs when knowledge requires a larger degree of justification. Rather, the skeptic claims that I have no degree of justification whatever for my claims about the external world. None. Let us call this view Full-Blooded Skepticism (Kornblith 2000, p. 25)

“In contrast, the High Standards Skeptic is perfectly willing to grant there are differences in degree of justification that people have for their various beliefs about the external world; he simply denies that we ever reach some very high standard required for knowledge” (Hannon 2016, p. 11). A Full-Blooded skeptic would ask: what justifies the proposition that we are not brains in vats? And she would answer that we have no such justification; how could we have such justification? As such, no standard, no matter how lenient, can possibly grant us any knowledge of the external world. As we see from Cohen, he mostly responds to the High Standards Skeptic, showing how we can have everyday knowledge even without being completely justified, by being “justified enough” in certain contexts. In this paper, in order to try to meet Cohen on his own terms as much as possible, skepticism has mostly been presented as the High Standards version. Though there is no reason why the Full-Blooded version is any less threatening (arguably, it is more threatening). Cohen's account of how knowledge ascriptions can be true by having a certain degree of justification is for nought if we never have any justification whatsoever for believing what we believe, yet this makes up the bulk of his theory. Any Full-Blooded Skeptic would be unfazed by this. The only way in which Cohen can combat the Full-Blooded skeptic is by reference to a certain intrinsic justification held by some claims (i.e. “no demon deceives us” or “the external world exists”). I already argued that such a thing, if it existed, would be too subjective. Lewis' account gives an answer to both kinds of skeptic, since he dismisses justification altogether; S may know p even if she has no justification for believing p. Moreover, Lewis of course avoids all the problems Cohen faces as a result of his reference to probability. I think these problems for Cohen are greater than the problems for Lewis. Indeed, I have no complaints about Lewis that I do not also have about Cohen. This is why I think that if Contextualism is to save us from skepticism, it must be by something resembling Lewis'

account more than Cohen. I do, however, have worries about the contextualist project as a whole.

Dust Under the Rug?

At this point in the paper, I will consider a few issues with contextualism in general. First, I will present a general objection to the contextualist project, and point out how, as much as it pains me to admit it, it is probably unsuccessful.

One thing to note about epistemic contextualism is that it does not give a response to any skeptical scenario. That is, the possibility that we may be brains in vats is made no less believable though any contextualist argument. Of course, no self-respecting epistemic contextualist is unaware of this. Cohen states that what his project tries to do is to solve a paradox (Cohen 1988, p. 177):

1. I know I have hands
2. I do not know that I am not a brain in a vat

These statements both seem true, but they also seem contradictory, thus we have a paradox. The paradox is solved by stating that the skeptic is right when she states “2”, and the anti-skeptic is also right in stating “1”. They are, however, right according to different standards. When the standards are held fixed, one is right and the other is wrong. However, we must ask ourselves, what good does this do? Surely, the skeptic's argument that we have no justification to believe what we believe is still of some concern? Haven't we just manipulated the requirements for what makes a knowledge ascription true without making our beliefs any more reliable? These are real issues, but not necessarily issues the contextualist find threatening to their account. After all, contextualists all agree that the skeptic is right once a skeptical scenario is considered, as they are ready to admit that they do not have any evidence to counter such scenarios. It should therefore be noted that the contextualist project is of limited use, even if it succeeds, as it does not ultimately try to refute any skeptics once they present their skepticism.

In order to illustrate this limitation, consider the following scenario: You are the coach of a basketball team. In basketball, taller players have an advantage over shorter ones, and suppose we consider a player “tall” just in case he is over 2 meters tall. All of your players are between 1.9 and 1.95 meters tall, and your team struggles because of a lack of tall players. Now, to solve this problem, you decide to change the meaning of “tall” to capture all players 1.9 meters tall or more. Now, all your players are tall, and you have solved the problem. Except, the team will not perform

any better, because what caused the issue was not that they failed to fall under the term “tall”. In order to solve the problem, you would need to have actually taller players in your team. While I think this analogy is not perfect, it serves to show that contextualism is a linguistic project, perhaps even more so than an epistemic one. It only addresses the terminology and semantics regarding the problem, solving the paradox in the process. But when someone is faced with a skeptical scenario (i.e. brains in vats), is the most frightening part about it that it renders the statement “I know I have hands” (among others) false? This worry is a linguistic one, as is somewhat analogous to someone learning that her coffee cup is not a cup, but rather a mug. This does not cause a massive amount of concern. Skepticism thus have a linguistic aspect, and an epistemic aspect. The linguistic aspect is the one that renders “I know I have hands” false, while the epistemic aspect is the one that undermines your grounds for believing you have hands. Contextualism only addresses the linguistic aspect of skepticism, while the epistemic aspect is left untouched. And arguably, the epistemic aspect is scarier. Generally, when non-philosophers consider skepticism, their worry is not the falsifying of knowledge ascriptions, but the fact that they have no way of determining what is true and what is not. This was Descartes' worry, and it has been the worry of skeptics before and since. In the face of this, the truth value of “S knows p” seems rather uninteresting. We can deem its truth conditions whatever we want, but the aspect of skepticism that has caused philosophers and non-philosophers concern for hundreds of years still remain. The contextualist's way of turning “S knows p” true again after the skeptic has challenged it is, then, in some way akin to shoving dust under the rug; for this reason, no skeptic will ever feel like epistemic contextualism addresses the actual point he is trying to make.

Now, contextualists may argue that it is in fact the other way around. It is not the contextualist that distorts the conception of knowledge to make knowledge attainable, but rather the skeptic that distorts it, making it unattainable. So a better analogy, according to contextualists, would be that a basketball coach changes the meaning of “tall” to refer to anyone at least one kilometer in height, making it impossible for any human to fall under the term “tall”. And so, the contextualist is not guilty of tampering with terminology, the skeptic is. The skeptic uses intuition pumps to make us question, and subsequently lose our knowledge. On this interpretation, contextualism is not about playing with words, but rather about knowledge, while the skeptic is guilty of that of which I just accused the contextualist.

At this point, it is the skeptic's word against that of the contextualist. As strongly as I feel the contextualist simply dodges the problem, I realise that it can just as easily be considered the other

way around. As Cohen puts it, “it is impossible for either side of the dispute not to beg the question against the other side” (Cohen 1988, p. 117). In the closing pages of this paper, I will discuss another way to consider why we say we know things, and at the same time say we do not. This view comes from Barry Stroud, and his book *The Significance of Philosophical Skepticism*.

Truth and Appropriateness

I still find myself discontent with the contextualist solution to the skeptical paradoxes. In this section I will explain, with some help from Barry Stroud, how it may be a better idea, and possibly also more true, to consider the skeptic right in all cases, not just some cases.

One argument for the contextualist conception of knowledge is that the utterance of skeptical scenarios often seem inappropriate and misplaced. For example, bringing up skeptical possibilities in a court of law, like dismissing an alibi because the witness possibly dreamt she was spending time with the defendant, is outrageous (Stroud 1984, ch.2, p. 9). It seems that we should grant that “the witness knows she spent time with the defendant” is true even though we cannot prove she was not dreaming, or at least this is how we seem to use language. But as Stroud points out, just because something is outrageous and inappropriate does not make it false. Claiming a witness does not know what she claims to know due to the dream argument is outrageous, but is it false? Stroud also brings up the following example.

For example, suppose I am at a party and my host asks me if I know whether my friend John, who was ill last week, will be coming to the party. I reply that I know he will be there, and when asked how I know I explain that he has now recovered, I have just talked to him on the telephone and he said he was coming right over; [...] Suppose that as I am leaving the party at the end of the evening and John has still not appeared my host turns on me and says ‘You should be more careful about what you claim to know. You said you knew John would be here and he isn't. [...] You didn't know any such thing!’ I think we find that response simply outrageous. But aside from the unreasonable abuse and insensitivity conveyed by it, the remark cannot be said to be totally false or without foundation. [...] What the cruel host says is an accurate description of my position (Stroud 1984, ch. 2, p. 16-17).

The point here is that even though I clearly did not know that John would come, it does not mean that I was unjustified in stating that I knew, and just because what the host said was inappropriate does not make it false. Stroud makes a distinction between conditions for knowledge and conditions for appropriate assertion of knowledge (ibid.). In the case of John and the party the assertion “I know John will come” was false, but appropriate. In other words, the conditions for truth was

unfulfilled, but the conditions for appropriate assertion was fulfilled. The host's rude remark “you didn't know” was the exact opposite; it fulfilled the conditions for truth, but not the conditions for appropriate assertion.

Thus far, things should be relatively uncontroversial; the question is, when the skeptic brings up his scenarios in everyday situations and concludes we do not know, does he violate the conditions for truth, or merely the conditions for appropriate assertion? It is my view, and apparently Stroud's view as well, that the skeptic does not violate the conditions for truth, in any context. Stroud also brings up another analogy, which I find useful:

Suppose that in wartime people must be trained to identify aircraft and they are given a quick, uncomplicated course on the distinguishing features of different planes and how to recognize them. They learn from their manuals, for example, that if a plane has features x, y, and w it is an E, and if it has x, y, and z it is an F. [...] Suppose that there are in fact some other airplanes, Gs say, which also have features x, y, z. The trainees were never told about them because it would have made the recognition of Fs too difficult; it is almost impossible to distinguish an F from a G from the ground (ibid., p. 22-23).

Stroud concludes that whenever an airplane-spotter discovers a plane with features x, y, z, and claims “I know it is an F”, he is speaking falsely, but appropriately and in accordance with his training. If told about the existence of the Gs, he would readily admit that he does not know it is an F. “He originally understood the word ‘know’ and applied that conception of knowledge fully reasonably and justifiably in particular cases, but (as we knew all along and he now would come to realize) he never knew [...] that the plane in the sky was an F” (ibid., p. 24). Notice the word “never”; on a contextualist interpretation, it may be claimed that he knew the plane was an F until he was told about the Gs. Alternatively, a contextualist would claim the situation not analogous, as the possibility of Gs would always be relevant. However, the manual was written without reference to Gs because it was considered not relevant. Bringing up the Gs simply does not help the war effort (ibid.), and so that possibility is ignored. However, he indeed never knew it was an F, not before or after the Gs were mentioned, and his conception of knowledge has not been changed.

To clarify, in the analogy, “I know it is an F” is an everyday knowledge attribution, “It may be a G” is a skeptical scenario, and the war effort is everyday life. The airplane-spotter's claim “I know it is an F” is false, it is still appropriate, and in accordance with the aims of the war effort. Bringing up Gs to not help matters. Likewise, everyday knowledge attributions can be false, while

still being appropriate and in accordance with the aims of everyday life. And suggesting evil scientists may be probing brains in vats does not help matters. Skeptical scenarios can undermine the truth of knowledge ascriptions without making it inappropriate to state them.

Why Choose Stroud's View?

This view has a few merits. First of all, it solves the skeptical paradox by claiming we do not know we have hands, avoiding contradiction. It also explains how the paradox shows up in the first place. The reason we think we know what we know is because there is a mix up between the conditions for truth of knowledge ascriptions and the conditions for asserting them appropriately. When a skeptic claims we know nothing because of a skeptical possibility, he is not speaking falsely or raising the standards for knowledge, merely speaking a truth that does not fulfill the conditions for appropriate assertion (in some cases it may be appropriate, but often it is not). And when we say we know things, we are not speaking truly, but appropriately, and in accordance with the aims of everyday life. This makes us think we know things, even though we do not.

Another merit of this view is that it explains why and how we can disregard skeptical possibilities in everyday life, in court, and other situations. If we accept contextualism as it is, it may seem as if skeptical possibilities must be considered once they are raised. As such, no court could ever convict a criminal if the defence brings up a skeptical scenario; after all, the *rule of attention* states that when a possibility is not ignored, it is not properly ignored, and thus it is within the domain of possibilities that must be eliminated. This is likely not an insurmountable problem for contextualism, but Stroud's view does a better job at explaining it. The reason we can discount skeptical possibilities in court is because they are outrageous with respect to the conditions for appropriate assertion, and considering them do not help the aims of the court. The claim "The witness does not know she spent time with the defendant, because she may have dreamt it" is true, and always is. It is however, merely outrageous to bring up in court. In this way, the context that contextualists speak of, does not affect the truth conditions of knowledge ascriptions, but rather the conditions for what is appropriate to assert.

While I do not have an infallible argument for why Stroud is right and Lewis is wrong, I find that Stroud tugs my intuition strings more so than Lewis, being more accommodating to the skeptic by giving the skeptic right at all times instead of just sometimes. But at the same time, Stroud is less accommodating to the skeptic in claiming he is quite often out of line when presenting his hypotheses, while Lewis sometimes may have trouble condemning the skeptic for bringing up

skeptic scenarios where they do not belong. In this way, Stroud is more accommodating to the skeptic on an epistemic basis, while he is less accommodating on a normative basis.

Conclusion

I have argued that the way in which we consider the difference between strict and lenient standards for the purposes of contextualism may be relevant to how well a particular brand of contextualism works. Cohen's probabilistic contextualism suffers from some trouble as a result of my conclusion that "S knows $P(p) \geq T$ " must be true if "S knows p" is to be true. I mentioned Sarah Moss as a potential ally, making "S knows $P(p) \geq T$ " plausibly true in some cases. However, Cohen's account still becomes subject to a particular form of the Regress Problem, as well as other minor issues, and thus I do not consider his project successful. David Lewis' salience contextualism is not subject to these problems, and thus I find his view more plausible. But even though Lewis' account may be consistent, I still find it leaves me discontent, as the skeptic issues has not truly been addressed. Until they are, I will find it hard to be convinced away from my skeptic tendencies. Barry Stroud's solution to these issues are more accommodating towards the skeptic, while still explaining why we need not (and should not) bring up skeptical possibilities at all times, and how we can reasonably assert everyday knowledge ascriptions even though they are most often (maybe even always) false.

References:

- Bostrom, N. (2001). "Are You Living in a Computer Simulation?" [online]. Available at: <https://www.simulation-argument.com/simulation.pdf> [Accessed 20 Oct. 2019].
- Cohen, S. (1988). "How to be a Fallibilist". *Philosophical Perspectives*, [online] Volume 2(Epistemology), pp. 91-123. Available at: https://www.jstor.org/stable/2214070?seq=1#metadata_info_tab_contents [Accessed 26 June 2019].
- Cohen, S. (1999). "Contextualism, Skepticism, and the Structure of Reasons". *Philosophical Perspectives*, [online] Volume 13(Epistemology), pp. 57-89. Available at: https://www.jstor.org/stable/2676096?seq=1#page_scan_tab_contents [Accessed 26 June 2019].
- DeRose, K. (2016). "Contextualism and Fallibilism" [online]. Available at: <https://cpb-us-w2.wpmucdn.com/campuspress.yale.edu/dist/c/1227/files/2016/03/Contextualism-and-Fallibilism-3-22-16-tod329.pdf> [Accessed 10 Nov. 2019].
- Descartes, R. (1641). "Meditations on First Philosophy". In: R. Ariew and E. Watkins, ed. (2009), *Modern Philosophy*, 2nd ed. Indianapolis: Hackett, pp. 35-68.
- Hannon, M. (2016). "Skepticism and Contextualism" [online]. Available at: Available at: <https://static1.squarespace.com/static/551587e0e4b0ce927f09707f/t/581f4a6bebbd1a1986459c39/1478445678167/Routledge.pdf> [Accessed 10 Nov. 2019].
- Klein, P. (1999). "Human Knowledge and the Infinite Regress of Reasons". In: E. Sosa, J. Kim, J. Fantl and M. McGrath, ed. (2008), *Epistemology*, 2nd ed. Oxford: Blackwell, pp. 165-185.
- Klein, P. (2003). "How a Pyrrhonian Skeptic Might Respond to Academic Skepticism". In: E. Sosa, J. Kim, J. Fantl and M. McGrath, ed. (2008), *Epistemology*, 2nd ed. Oxford: Blackwell, pp. 35-50.
- Kornblith, H. (2000). "The Contextualist Evasion of Epistemology". *Philosophical Issues*, [online] Volume 10(Skepticism), pp. 24-32. Available at: <https://www.jstor.org/stable/3050560?>

[seq=1#metadata_info_tab_contents](#) [Accessed 20 Oct. 2019].

Lewis, D. (1996). “Elusive Knowledge”. *Australasian Journal of Philosophy*, [online] Volume 74(No.), pp. 549-567. Available at: <https://fitelson.org/epistemology/lewis.pdf> [Accessed 14 Oct. 2019].

Moss, S. (2018). *Probabilistic Knowledge*. [online]. Available at: <https://www.oxfordscholarship.com/view/10.1093/oso/9780198792154.001.0001/oso-9780198792154> [Accessed 15 Sep. 2019].

Stroud, B. (1984). *The Significance of Philosophical Skepticism*. [online]. Available at: <https://www.oxfordscholarship.com/view/10.1093/0198247613.001.0001/acprof-9780198247616> [Accessed 22 Oct. 2019].

Williamson, T. (2013). “Very Improbable Knowing”. [online]. Available at http://media.philosophy.ox.ac.uk/assets/pdf_file/0015/19302/veryimprobable.pdf [Accessed 10 Nov. 2019].