

Epistemic invariantism and contextualist intuitions

Alexander Dinges

alexander.dinges@uni-hamburg.de

(forthcoming in *Episteme: A Journal of Individual and Social Epistemology*)

Abstract: Epistemic invariantism, or invariantism for short, is the position that the proposition expressed by knowledge sentences does not vary with the epistemic standard of the context in which these sentences can be used. At least one of the major challenges for invariantism is to explain our intuitions about scenarios such as the so-called bank cases. These cases elicit intuitions to the effect that the truth-value of knowledge sentences varies with the epistemic standard of the context in which these sentences can be used. In this paper, I will defend invariantism against this challenge by advocating the following, somewhat deflationary account of the bank case intuitions: Readers of the bank cases assign different truth-values to the knowledge claims in the bank cases because they interpret these scenarios such that the epistemic position of the subject in question differs between the high and the low standards case. To substantiate this account, I will argue, first, that the bank cases are underspecified even with respect to features that should uncontroversially be relevant for the epistemic position of the subject in question. Second, I will argue that readers of the bank cases will fill in these features differently in the low and the high standards case. In particular, I will argue that there is a variety of reasons to think that the fact that an error-possibility is mentioned in the high standards case will lead readers to assume that this error-possibility is supposed to be likely in the high standards case.

1. Introduction

Epistemic invariantism, or invariantism for short, is the position that the proposition expressed by knowledge sentences does not vary with the epistemic standard of the context in which these sentences can be used. At least one of the major challenges for invariantism is to explain our intuitions about scenarios such as the so-called *bank cases*. These cases

elicit intuitions to the effect that the truth-value of knowledge sentences varies with the epistemic standard of the context in which these sentences can be used. In this paper, I will defend invariantism against this challenge by advocating the following, somewhat deflationary account of the bank case intuitions: Readers of the bank cases assign different truth-values to the knowledge claims in the bank cases because they interpret these scenarios such that the epistemic position of the subject in question differs between the high and the low standards case. To substantiate this account, I will argue, first, that the bank cases are underspecified even with respect to features that should uncontroversially be relevant for the epistemic position of the subject in question. Second, I will argue that readers of the bank cases will fill in these features differently in the low and the high standards case. In particular, I will argue that there is a variety of reasons to think that the fact that an error-possibility is mentioned in the high standards case will lead readers to assume that this error-possibility is supposed to be likely in the high standards case. The structure of the paper is as follows: In the first section, I will present the challenge from the bank cases for invariantism in more detail. In the second section, I will lay out the just indicated invariantist response to this challenge. In the last section, I will show how to deal with some objections to my view.

2. The Challenge for Invariantism

Consider the following scenarios, the so-called *bank cases*:

Low standards: Hannah and Sarah are driving home from work on a Friday afternoon. They plan to stop at the bank to deposit their paychecks, but as they drive past the bank they notice that the lines inside are very long. Although they generally like to deposit their paychecks as soon as possible, it is not especially important in this case that they be deposited right away. Hannah calls Bill and asks whether the bank will be open on Saturday. Bill tells her that the bank will be open and that he remembers driving by last Saturday seeing that it was open until noon. Hannah says, “Bill told me that the bank will be open. Let’s drive straight home and return to deposit the paychecks tomorrow morning.” Sarah responds, “Maybe the bank won’t be open tomorrow. Lots of banks are closed on Saturdays. Does Bill know the bank will be open tomorrow?” Hannah replies,

“Bill knows the bank will be open tomorrow.” It turns out that the bank is open on Saturday.

High standards: Hannah and Sarah are driving home from work on a Friday afternoon. They plan to stop at the bank to deposit their paychecks, but as they drive past the bank they notice that the lines inside are very long. Hannah and Sarah have just written a very large check, and if the money from their pay is not deposited by Monday, it will bounce, leaving them in a very bad situation with their creditors. And, of course, the bank is not open on Sunday. Hannah calls Bill and asks whether the bank will be open on Saturday. Bill tells her that the bank will be open and that he remembers driving by last Saturday seeing that it was open until noon. Hannah says, “Bill told me that the bank will be open. Let’s drive straight home and return to deposit the paychecks tomorrow morning.” Sarah reminds Hannah of how important it is to deposit the check before Monday and says, “Banks are typically closed on Saturday. Maybe this bank won’t be open tomorrow either. Banks can always change their hours. Does Bill know the bank will be open tomorrow?” Hannah replies, “Well, no, Bill doesn’t know the bank will be open tomorrow.” It turns out that the bank is open on Saturday.

What should we say about the truth-value of the knowledge claims in the low and the high standards case? Many authors have endorsed the following intuitions: On the one hand, it seems true for Hannah in the low standards case to say, “Bill knows the bank will be open.” On the other hand, it seems true for Hannah in the high standards case to say, “Bill doesn’t know the bank will be open.” Let’s refer to these data as *the Contextualist Data*.¹

Opponents of invariantism use the Contextualist Data to argue against invariantism in the following way: Bill’s epistemic position is the same in the low and the high standards case. As DeRose puts it,

¹ The Contextualist Data may be doubted, but since they have been endorsed by many philosophers, I will just grant them for the purposes of this paper. Besides, closely related data have been empirically confirmed e.g. in (Schaffer and Knobe 2012), (Hansen and Chemla 2013), (Nagel et al. 2013) and (Buckwalter 2014). See below for some further remarks on the plausibility of the Contextualist Data.

our subject's [Bill's] situation is identical in our two third-person cases except for the difference that, far away from her, different conversations, which she is and will remain oblivious to, are taking place about her. (DeRose 2009: 64)

Correspondingly, Bill should either know that the bank will be open in both cases or fail to know that the bank will be open in both cases. So, assuming invariantism, it should either be wrong for Hannah to say, "Bill knows the bank will be open," or it should be wrong for her to say, "Bill doesn't know the bank will be open." Contrary to these predictions, however, the Contextualist Data have it that both these sentences seem true in their respective scenario. Thus, invariantists are unable to explain the Contextualist Data.

Invariantists have tried to respond to this challenge in various ways, but none of their proposals is without problems.² Instead of rehearsing these proposals here, though, I will

² *Pragmatic invariantists* have tried to explain the Contextualist Data based on alleged conversational implicatures of Hannah's knowledge claims. See e.g. (Rysiew 2001). Blome-Tillmann (2013) points out a range of problems for such views. *Subject-sensitive invariantists* have tried to explain the Contextualist Data by arguing that one's epistemic position may be affected by seemingly non-epistemic facts, such as how much is at stake. See e.g. (Stanley 2005). As DeRose points out in the just quoted passage, these proposals fail once we consider the above "third-person" bank cases, where the relevant speaker does not ascribe knowledge to herself but to a third party not participating in the relevant conversation. Analogous considerations hold for the *belief removal theory* as put forward e.g. in (Bach 2005), according to which the subject in the high standards case stops believing that the bank will be open in the face of the heightened stakes and the mentioned error-possibility. *Psychological error theorists* have tried to explain the Contextualist Data based on the idea that, for one or the other psychological reason, we hold mistaken beliefs about Bill's epistemic position in the bank cases. See e.g. (Williamson 2005). Nagel (2010) presents serious objections to at least the most prominent versions of this kind of position (those that appeal to the so-called "availability heuristic"). The view I will put forward below is straightforwardly compatible with the rejection of the first three theories. However, it bears close similarities to psychological error theories. First, it also appeals to the psychology of the readers of the bank cases. Second, as most versions of the psychological error-theory, it trades on how likely we consider a given error-possibility to be. Still, the psychological mechanisms I will employ differ from the mechanisms psychological error-theorists have employed (for example, my account does not refer to the availability heuristic, see FN 9). Besides, it is at least unclear whether my view entails that the readers of the bank cases are mistaken in any sense. (To this extent, it is unclear whether my view genuinely *is* a form of psychological error-theory rather than just being similar to it. See FN 5.) My view also bears similarities to subject-sensitive invariantism and the belief removal theory in that it also holds that the relevant subject's

simply put forward a novel, somewhat deflationary account of the Contextualist Data. In particular, I will argue, first, that the descriptions of the bank cases leave open epistemically relevant aspects of the situation that need to be filled in before we can assess Hannah's knowledge claims. Second, I will show that these aspects will be filled in such that Bill's epistemic position does differ between the bank cases. Given this result, the invariantist is at liberty to say that Bill's epistemic position is good enough for knowledge only in the low standards case. Thus, she can explain the Contextualist Data by simply pointing out that, with respect to the bank cases as they are interpreted by their readers, the relevant judgments are perfectly correct.^{3, 4, 5}

epistemic position differs between the bank cases (at least once we take into account how the readers of the bank cases interpret these scenarios). It diverges from these positions in that this difference is neither supposed to be due to seemingly non-epistemic features of the scenarios nor due to the subject's beliefs. Besides, my proposal does apply to the third-person cases, as the subsequent discussion will show.

³ It may be interesting to note that DeRose himself employs a strategy along the above lines to undermine empirical findings purportedly supporting contrastivism (an anti-invariantist view defended e.g. in (Schaffer 2004)). In particular, he also argues that "the participants in the two groups [assessing the relevant scenarios] are construing the situation in relevantly different ways." (2011: 109)

⁴ Couldn't we block any such account by simply stipulating that Bill's epistemic position is the same in the bank cases? I will address this worry in the objections and replies section below.

⁵ Is the proposed view an error-theory? This is somewhat hard to say. Error-theorists typically maintain that readers of the bank cases hold mistaken beliefs about one or the other aspect of the bank cases. The relevant beliefs in my account, however, are about aspects of the bank cases the descriptions of these cases leave unspecified. This does not mean that there is no sense in which these beliefs could be mistaken. For example, we could say that they are mistaken in that they go against the *intended* interpretation of the bank cases. Similarly, they could be mistaken in that they are *irrational* interpretations of these scenarios. I think that the subsequent discussion shows that the suggested interpretation is not at all irrational and that, even if it is unintended, the relevant intentions cannot be expected to be transparent to the readers of the scenarios. But be that as it may. I am happy to grant that my view is some sort of error-theory. All I seek to establish is that the mistake involved is at best a very mild one that can robustly be explained in terms of the considerations to come.

3. Responding to the Challenge

To begin with, note that nothing is said explicitly in the descriptions of the bank cases about how likely it is supposed to be that Hannah and Sarah's bank changes its hours. Nothing is said, for example, about how frequently it changed its hours before, how frequently shops in general change their hours in Hannah and Sarah's environment, how reliable the director of the bank is supposed to be, etc. Nothing is said about all these aspects of the situation. Still, quite plausibly, they affect Bill's epistemic position with respect to the claim that the bank will be open. For, the more likely it is that the bank changes its hours, the easier it will be for Bill to be wrong about the opening hours of the bank. And, plausibly, the easier it is to be mistaken, the weaker one's epistemic position. Thus, readers of the bank cases will have to fill in these details in one way or another before they can assess whether Hannah's knowledge claims are true or false.⁶

But how will they fill in these details? In what follows, I will present three mutually supportive, yet self-standing arguments for the claim that the details will be filled in differently in the different bank cases. In particular, I will argue that the details will be filled in such that the possibility that the bank has changed its hours is less likely in the low than in the high standards bank case. According to the first argument, this is so because we are reminded of this possibility when we read the high standards scenario. According to the second argument, this is so because Sarah implicates that this possibility is likely when she mentions it. According to the third argument, this is so because this possibility would not normally occur to Sarah unless it was likely. I take the first argument to be the most original contribution to the debate. The second and the third argument draw heavily upon proposals already present in the literature. However, I will put a novel spin on these proposals by embedding them into the argumentative framework above, according to which

⁶ Note that, when I speak of likelihoods here and in what follows, I have in mind some kind of more or less objective sort of likelihood that is relatively independent of the evidence available to particular subjects. As just indicated, for example, the likelihood of the possibility that Bill is mistaken is supposed to depend on how frequently the bank has changed its hours before, whether or not Bill is aware of how often the bank has changed its hours before. DeRose (2009: 22f) and Gerken (2011: 56), for instance, explicitly endorse the idea that a subject's epistemic position depends on the likelihood of error-possibilities in this objective sense.

the Contextualist Data must be explained by appeal to how readers of the bank cases fill in these scenarios.

Before we can spell out the just mentioned arguments in more detail, we first have to ask how readers of a scenario generally fill in the missing details of a scenario. This is ultimately an empirical question, but the following seems like a principle that should guide us at least to a certain extent:

Normalcy Projection (NP): Subjects will tend to fill in the missing details of a scenario in accordance with what they think is normal.

The idea behind this principle is that, if the author of a scenario fails to specify a relevant bit of the scenario, then, presumably, this part of the scenario is not supposed to be special in any way. Hence, we just fill in what we consider normal.

(In what follows, I will sometimes say that this-and-that *will be assumed to be* thus-and-so in the bank cases instead of using the cumbersome phrase that the readers of the bank cases will fill in these cases such that this-and-that is thus-and-so. For example, I will say that Bill's epistemic position *will be assumed to be* weaker in the high than in the low standards case. This is a mere terminological simplification, though, with no substantive import.)

The Argument from Reminding

Here then is a detailed version of the first argument to the conclusion that the bank will be assumed to be more likely to change its hours in the high standards case. Plausibly, when we don't think of an error-possibility, we tend to act as if we thought that its likelihood is zero even if we already know (in some sort of unconscious or tacit way) that the likelihood is not zero. For example, suppose that you know full well that robins and bluebirds are equally likely to be seen in your area. Suppose further that you see an orange-breasted bird that, for all you see, is either a robin or a bluebird. If you don't think of the possibility that the bird is a bluebird (maybe because it's been a while since you last saw a bluebird), you will presumably still say, "That's a robin!" This is so because you treat the likelihood of

unthought-of possibilities as zero.⁷ When subjects read the low standards case, they similarly don't think of the possibility that (normal) banks change their hours. Hence, given the previous observation, they will treat the likelihood of this possibility as zero. Given (NP), they will project this likelihood into the scenario. Thus, they will assume that the likelihood of the possibility that Hannah and Sarah's bank changes its hours is zero in the low standards case.

In the high standards case, the situation is different. It seems clear that once we are reminded of a known-to-be-likely possibility, we no longer treat its likelihood as zero even if we haven't been taking it into account before. For example, when I remind you, "It could also be a bluebird," you'll immediately retract your original assertion of "That's a robin!" because you remember that this possibility is likely.⁸ Similarly, when subjects read the high standards case, they are reminded of the possibility that (normal) banks change their hours because this possibility is mentioned in the scenarios. Hence, they also realize that its likelihood is, though maybe low, certainly not zero. Again, given (NP), they will project this assumption into the scenario and thus assume that it is likely at least to a certain degree that Hannah and Sarah's bank changes its hours. So, the assumed likelihood of this possibility will be higher in the high than in the low standards case.⁹

⁷ I take it that the above principle is fairly obvious and hence need not be backed up by empirical results. Still, of course, it is in the end an empirical psychological principle. For those who thus seek empirical corroboration, I refer to (Gerken 2011: 50; 2012: 155). The right-to-left direction of what he calls the "Principle of Contextual Salience" seems to correspond at least roughly to the above thesis. Gerken discusses this thesis in a psychologically much more informed fashion than I am able to do.

⁸ The above claim is closely related to the left-to-right direction of Gerken's Principle of Contextual Salience, which has it that "[n]ormally, if for an agent, *A*, *q* is a contextually salient alternative to *S*'s knowledge that *p*, then *A* processes *q* as an epistemically relevant alternative to *S*'s knowledge that *p*." However, my contention here is much more moderate and therefore hopefully uncontroversial. It roughly says that normally, if for an agent, *A*, an alternative, *q*, is salient, and *A* knows that *q* is likely to a certain degree, then *A* "processes" *q* as likely to that degree (which, in turn, may mean that *A* processes *q* as epistemically relevant, depending on how the notion of epistemic relevance is understood and depending on how likely *q* is supposed to be).

⁹ Note that the just described psychological mechanism is very different from the mechanism that psychological error-theorists have typically employed to explain the Contextualist Data. Psychological error-

The Argument from Implicatures

Here is a detailed version of the second argument for this conclusion. It is based on the idea that, first, Sarah implicates that it is likely to a certain degree that the bank changes its hours and, second, that this implicature will be assumed to be true in the high standards case given (NP). Let's discuss these claims in turn.

Why should Sarah implicate that it is likely to a certain degree that the bank changes its hours? Here, we can draw on the work of Dougherty and Rysiew (2009). They pursue a project that is closely related to mine, namely, they try to explain, from an invariantist (intellectualist and fallibilist) point of view, why so-called “concessive knowledge attributions” sound odd (a concessive knowledge attribution being a claim of the form “I know that p , but it is possible that q ” where q entails not- p). We need not bother about the details of their account here. All that is relevant is that, in the course of their discussion, they plausibly argue for the following claim: “In saying that p is possible one implies that there is some real chance that p .” (2009: 130) (Dougherty and Rysiew do not say much about what a “real” chance is supposed to be. Still, they assume that, if there is a real chance of error, then this “may well be enough to prevent one from knowing.”) Following Dougherty and Rysiew, the above implicature can be derived as follows:¹⁰ The mere existence of an error-possibility will not be relevant for the question of whether someone knows that a certain proposition is true. (Otherwise, we would presumably be wedded to a fairly strong form of skepticism.) So, when Sarah says, “Banks can always change their hours,” she will say something irrelevant to the extent that all she says is that this possibility exists. The irrelevance of her claim will thus trigger the search for an

theorists usually argue that we *overestimate* the likelihood of salient error-possibilities. The proposal above rather has it that, in some sense, we *underestimate* the likelihood of *non*-salient error-possibilities. My proposal is thus related to the “shallow processing” account discussed in (Schaffer and Knobe 2012: 26). Schaffer and Knobe argue that such accounts lead to skepticism. I will address this objection below.

¹⁰ See (Dougherty and Rysiew 2009: 128f). See (Bach 2008: 79) for a similar reasoning.

implicature. And this implicature plausibly is that the possibility that banks change their hours not only exists, but also is likely to a relevant degree.^{11, 12}

Why should readers of the bank cases assume that this implicature is supposed to be true? We do not normally try to mislead people. So, given (NP), Sarah will have to be assumed to believe what she implicates (that there is a real chance that the bank changes its hours). Now, in the absence of misleading evidence, we normally hold true beliefs about our everyday concerns (such as the opening hours of our bank). So, given that no reason is given in the bank cases to think that Sarah should have been misled in any way, and given (NP), Sarah's belief will be considered true. Thus, we have a second argument for the conclusion that the possibility that the bank changes its hours will be assumed to be more likely in the high than in the low standards case.¹³

The Argument from Normal Cognitive Processes

Here is a detailed version of the third argument for this conclusion. As Bach plausibly points out, if “our cognitive processes are operating well, generally the thought of a possibility contrary to something we're inclined to believe occurs to us only if it's a realistic possibility, not a far-fetched one” (Bach 2010: 121). Thus, the possibility that the

¹¹ The derivation of this implicature relies on assumptions about the semantics of possibility claims that may not be entirely uncontroversial (in particular, that their semantic content by itself is irrelevant for the question of whether someone knows a given proposition). Addressing this issue, however, goes beyond the scope of the present project. See (Dougherty and Rysiew 2009) for further discussion.

¹² I tacitly made the transition from the idea that a given possibility is “real” to the idea that the possibility is likely. It may be doubted, however, whether this transition is legitimate. For example, one may think that whether a possibility is real depends on how much is at stake. Thus, one may hold, a possibility can be real even if it is unlikely, namely, when much is at stake. (See e.g. (Fantl and McGrath 2009: 22f) for such a conception of the realness of a possibility.) Here is my response to this worry. Even those who want to say that the realness of a possibility depends on how much is at stake do not want to say that the realness of a possibility depends *only* on what is at stake. Instead, they will say that a possibility is real if it is likely enough for the stakes at hand (or something along these lines). Thus, if a possibility is real, it still follows that it is likely at least to a certain degree. And this is all that is required for the above argument to go through.

¹³ For a suggestion that goes in in the direction of this second argument, see May's comments on <http://certaindoubts.com/?p=1373>.

bank changed its hours should normally occur to Sarah only if this possibility is likely to a certain degree. So, given that this possibility does occur to Sarah in the high standards case, we have to assume that it is likely there in order to preserve normalcy. Hence, given (NP), we have a further argument for the claim that the possibility that the bank changes its hours will be assumed to be less likely in the low than in the high standards bank case.¹⁴

Note that the presently described mechanism is importantly different from the implicature mechanism above. For example, the implicature mechanism arguably won't apply when we consider bank cases where the error-possibility is not mentioned in the high standards case, but only considered in thought by the knowledge-ascribing subject.¹⁵ After all, thoughts (as opposed to mentionings) plausibly don't have implicatures. The present mechanism, however, straightforwardly does apply to these cases, for, to repeat, an error-possibility normally occurs to somebody (in thought or in speech) only if this possibility is likely.

To sum up, we have seen that there are three straightforward reasons to hold that the possibility that the bank changes its hours will be assumed to be less likely in the low than in the high standards bank case. Given that, the invariantist is free to say that Bill's epistemic position is good enough for knowledge only in the low standards case. Thus, she has a straightforward explanation of the Contextualist Data: the manifested intuitions are correct with regard to the bank cases as they are interpreted by their readers.¹⁶

¹⁴ The points from FN 12 apply here as well.

¹⁵ See e.g. (Baumann 2011) for such bank cases.

¹⁶ Let me note at this stage that I am not entirely certain to what extent we should respect the Contextualist Data. In particular, when we look at the results from the studies on the bank cases mentioned above, we see that, even though we do seem to evaluate the low and the high standards case differently, the difference is rather small. A more plausible account of these results might be the following: Bill is a borderline case of a knower in the low and the high standards case. However, given the just described difference in the likelihood of the relevant error-possibility, he is somewhat closer to being a knower in the low rather than the high standards case; thus the (small) difference in our evaluations. This modification of my account would also answer the potential worry that "it is unclear exactly why a slight increase in the apparent odds of danger in cases like [the high standards case] would be sufficient to overturn the positive assessment of cases like [the low standards case]." (Nagel 2010: 299) The difference may not suffice to overturn this assessment. But

4. Objections and Replies

In what follows, I will respond to a range of potential objections to the view just proposed. It will be useful to have a label for this view, so let's call it *projective invariantism* (because it trades on the idea that readers of the bank cases *project* certain assumptions into the scenarios).

Skepticism

Here is the first objection to projective invariantism. Schaffer and Knobe (2012: 27) object to a similar position that it would commit the invariantist to “the quasi-skeptical view that in fact Hannah [in our case, Bill] does not know that the bank will be open on Saturday” and suggest “that most invariantists will already reject the strategy at this point.”¹⁷ Since projective invariantism also entails that Bill doesn't know that the bank will be open in the high standards case, they would presumably level the same objection against this view. The reasoning could go as follows: In ordinary conversations, we generally ascribe knowledge to subjects whose epistemic position is the same as, or even weaker than, Bill's in the high standards case. So, saying that Bill fails to know that the bank will be open in the high standards case commits one to the thesis that our ordinary practice of ascribing knowledge is generally mistaken. This result, however, looks like a form of skepticism that almost everybody wants to avoid.

To respond, I think that there is just no basis for the assumption that, in general, we ascribe knowledge to subjects whose epistemic position is the same as, or even weaker than, Bill's in the high standards case. The projective invariantist, at least, has reasons to reject this assumption: According to the Contextualist Data, we would intuitively judge that it is correct for Hannah in the high standards case to say, “Bill doesn't know that the bank will be open.” Thus, assuming invariantism, we have a straightforward reason to hold that,

since, according to the empirical results, our judgments aren't overturned to begin with (they only change slightly), this is not a problem for my view.

¹⁷ Gerken (2011: 54f) also endorses an objection along these lines. This may be part of the reason for why he doesn't apply his Principle of Contextual Salience (see FN 7 and 8) in the way I propose to apply it.

in general, we wouldn't ascribe knowledge to subjects whose epistemic position is the same as, or even weaker than, Bill's in the high standards case.

Of course, according to the Contextualist Data, we would also intuitively judge that it is correct for Hannah in the low standards case to say, "Bill knows that the bank will be open." However, for the projective invariantist, this will show at best that we would ascribe knowledge to subjects whose epistemic position is the same as Bill's *in the low standards case*. And this result is compatible with the result from the previous paragraph that we wouldn't ascribe knowledge to subjects whose epistemic position is the same as Bill's *in the high standards case*.

Making the Likelihood Explicit

Here is a second potential objection to projective invariantism. Stanley alludes to this objection when he writes that "even when we have full knowledge of the likelihood of various counter-possibilities, it does not affect our judgments." (Stanley 2005: 101) Applied to projective invariantism, the objection could be developed as follows: Projective invariantism crucially relies on the idea that the bank cases are underspecified when it comes to how likely it is that the bank changes its hours. For this allows readers to fill in the low and the high standards case in different ways. Thus, if projective invariantism holds, the Contextualist Data should disappear once we explicitly stipulate how likely it is that the bank changes its hours in the bank cases. To make the Contextualist Data disappear, it should suffice, for example, to end the high standards case with a claim to the effect that the likelihood of the possibility that the bank changes its hours is zero. Stanley's quote above can be seen as the objection that this prediction is mistaken.

To respond, I think that Stanley is just wrong about this latter judgment. At the very least, the judgment is unfounded. It seems quite plausible that the Contextualist Data would disappear if the scenarios were specified in the mentioned respect. Suppose, for example, that we end the high standards case with claims along the following lines:

Note that, even though Sarah mentions the possibility that the bank has changed its hours, this possibility is utterly unlikely. The bank hasn't changed its hours before. Banks in Hannah and Sarah's environment are generally very reliable. And even if the

bank had changed its hours, this would have been properly announced at least two weeks before, so Bill would most likely have known.

It seems not at all clear that, once we add the above paragraph, people would still agree that Bill fails to know that the bank will be open in the high standards case.¹⁸

Making Sameness Explicit

Here is a slight variation on the previous objection. Projective invariantism no longer applies once we explicitly stipulate that Bill's epistemic position is the same in the low and the high standards case after having described these cases. (Similarly, the view no longer applies when we consider a single scenario containing the dialogues from both the low and

¹⁸ A referee has suggested that, even if we add the above paragraph, our judgments about the bank cases may still vary if it's built into the high standards case that the subject should play it safe and wait in the lines today and built into the low standards case that the subject is fine to come back tomorrow. Thus, our judgments about the knowledge claims in the bank cases would be sensitive to whether or not it is fine for the subject to act as if the putatively known proposition is correct. (To properly construct such cases, we would presumably have to consider first-person versions of the bank cases, in which the subject faced with the decision to wait in the lines or come back tomorrow is the same as the subject of the knowledge claims.) Now, I am somewhat doubtful that such cases would indeed elicit the relevant intuitions. After all, the present high standards case would plausibly just trigger the puzzling question of why the subject should not come back tomorrow given that it is so unlikely that the bank changes its hours. But even if they do elicit the relevant intuitions, it may be possible to explain these intuitions in the spirit of the view put forth in this paper: Once we state in the high standards case that the subject should wait in the lines, we suggest that there is *some* likely possibility preventing the bank from opening tomorrow. Of course, it is now explicitly stated that this is not the possibility that the bank changes its hours. But there are alternatives. For example, it may be typical for banks in Hannah and Sarah's environment to open on Saturdays only once a month. Thus, again, readers of the bank cases may assume that the subject's epistemic position is supposed to be weaker in the high standards case. This account would have to be worked out in more detail. But even if it fails (and even if the relevant cases trigger the relevant intuitions in the first place), I don't think this would be a devastating result. As long as I can explain the effect of salient alternatives, this seems progress enough. An account of the connection between knowledge and actionability may just have to wait for another occasion.

the high standards case.¹⁹) Even so, the Contextualist Data would not disappear. Thus, projective invariantism fails to explain all relevant data.

To see that this objection fails, we have to get clear on what exactly the Contextualist Data amount to. I grant that, even when the bank cases are presented in the just proposed way, the following pattern would emerge: While reading the low standards case (before we arrive at the high standards case and the stipulation that Bill's epistemic position is the same), we would judge that the knowledge ascription in the low standards case is true. Once we have read the high standards case and the stipulation, we would judge that the knowledge denial in the high standards case is true. This pattern, I think, is driving the Contextualist Data. And projective invariantism straightforwardly explains it: The projective invariantist can simply argue that we change our interpretation of the scenarios once we are reminded of the error-possibility in the high standards case.

Projective invariantism would be in trouble only if, even after having read the high standards case and the stipulation, we would still say that the knowledge ascription in the low standards case is true. It is very much unclear, however, whether this would be so. Even DeRose (2009: 2) admits that to obtain the Contextualist Data we should consider the low standards case "without also thinking about" the high standards case (and *vice versa*). The point can be strengthened by the observation that most people would grant so-called "retraction data."²⁰ Most people would agree that, once epistemic standards go up, we do not only say that we don't know this or that proposition, we also reject our earlier claims to the effect that we do know the proposition. In the same way, it seems very plausible that, once we have read the high standards scenario, we will no longer say that the knowledge ascription in the low standards case is true.

Far-Fetched Alternatives

Here is another closely related objection. One could argue that we would obtain the Contextualist Data even if we let Sarah mention an alternative that we don't consider likely at all. For example, one could argue that the Contextualist Data would still be present if

¹⁹ See e.g. (DeRose 2009: 4f) for such a scenario.

²⁰ See e.g. (MacFarlane 2005).

Sarah mentions the possibility, say, that a meteor destroys the bank before Saturday. This result would be problematic because projective invariantists explain the Contextualist Data at least partly on the grounds that we do consider the alternative in the high standards case likely (and are reminded of that once the possibility is mentioned).

Again, my response is just to deny the alleged data. In fact, I think the real data favor my proposal. Gerken seconds this response in the following passage:

[A]n important consideration concerning very far-fetched salient alternatives cases must be noted. Consider, for example, the alternative that the matter in *S*'s car has spontaneously reorganized in the form of a giant lizard (MacFarlane 2005). This salient alternative to *S*'s knowledge that the car is in the driveway does not generate strong intuitions to the effect that *S* does not know that the car is in the driveway (MacFarlane 2005). [FN: We have at present no empirical evidence about far-fetched salient alternatives. But the intuitive judgment appears to be widely agreed upon.] The absence of an intuition in this case—or, at least, the asymmetry in strength of intuition—needs to be explained as much as the presence of intuitions in cases of less far-fetched salient alternatives. (Gerken 2012: 141f)

If Gerken is correct, it is indeed a good thing that projective invariantism does not predict contextualist effects when very far-fetched alternatives are mentioned; there simply are no such effects.²¹

Skeptical Arguments

The following is a further objection in the vicinity of the previous considerations. Suppose a skeptic, instead of merely making an error-possibility salient, presents arguments to the conclusion that we don't know that the error-possibility does not obtain and then appeals to epistemic closure to reach the skeptical conclusion that we don't know what we thought we

²¹ Admittedly, Gerken does not directly say that the Contextualist Data are absent when very far-fetched alternatives are mentioned. He leaves open the possibility that the data should only be weaker. But this is unproblematic for the projective invariantist, for at least some of the mechanisms described above to accommodate the Contextualist Data may apply even when the mentioned error-possibility is very far-fetched. In particular, readers may still take Sarah's claim to have the implicature described, or assume that the possibility she mentions is supposed to be likely in the bank cases in order to make sense of why it occurs to Sarah in these cases.

know. Suppose, for example, that the skeptic points out that we don't know that we are not brains in a vat because our sensory experiences would be exactly the same as they are if we were brains in a vat. Suppose further that she now appeals to one or the other epistemic closure principle to conclude that we don't even know we have hands. One might be inclined to think that people tend to respond to such arguments by granting that they don't know they have hands. And it might further be held that projective invariantism is unsuited to explain this datum. In support of this latter contention, Nagel, for example, argues that, when

we succumb to the temptation of rescinding knowledge ascriptions under sceptical pressure, we can at least ostensibly continue to judge the odds of being in a sceptical scenario as very low: the sceptic might move you into suspension of judgement even after granting that the odds of being a BIV are one in a quadrillion, if he is skilled at getting you to appreciate that things would appear *exactly* the same to you under that condition. (Nagel 2010: 299)

I fully agree with Nagel that projective invariantism cannot account for the fact (if it is fact) that we retract knowledge claims in the face of skeptical arguments, at least when these arguments involve far-fetched alternatives such as that of being a brain in a vat. For, as indicated by Nagel, none of the mechanisms posited above to account for the Contextualist Data applies to these reactions. First, we will not be reminded that it is likely to a relevant degree that we are brains in vats when this possibility is pointed out to us. After all, we presumably consider this alternative extremely unlikely. Relatedly, the skeptic may happily cancel any implicature to the effect that this possibility is likely without undermining her case. Finally, it seems clear that the skeptical error-possibility comes to the mind of the skeptic not because it is likely but because the skeptic has been engaging in philosophy.

But how problematic is the result that projective invariantism cannot explain our assumed reactions to skeptical arguments? I don't think it is very problematic. After all, it seems to be very much unclear that we do indeed tend to retract knowledge claims in the face of skeptical arguments. Many, if not most people are entirely unmoved by such considerations.²² What does seem correct is only that *some* people become convinced by

²² See (Feldman 2001: 77), (Bach 2010: 108f) and (Davis 2007: 436) for related observations.

skeptical arguments (some philosophers, for example) and that, for most people, it is difficult to see where exactly skeptical arguments go awry (even if most people think that these arguments just *have to* go awry at some point or other). Admittedly, projective invariantism cannot explain these residual data either. But this result is problematic only if there are no plausible accounts of these data the projective invariantism could adopt. This is not the place to argue that there are such accounts. It seems clear, however, that the burden of proof is on those who want to say that there are none. For example, the projective invariantist could argue that skeptical arguments may seem convincing because they involve one or the other non-obvious argumentative fallacy.²³

Mere Stakes Effects

The next and final worry has it that we would obtain the Contextualist Data even if we only changed the stakes between the low and the high standards case and not the mentioned error-possibilities.²⁴ This might again seem to undermine my proposal, for if no error-possibility is mentioned, we arguably won't be reminded of an error-possibility, the relevant implicature will also not be present, and there will be no reason to think that an error-possibility occurred to any of the subjects involved in the bank cases.

Here are my responses to this objection. First, it is far from clear that mere stakes effects cannot be explained on the basis of projective invariantism. Drawing on results from psychology, Buckwalter and Schaffer (2013: 22) argue that high stakes “can be expected to trigger anxiety in the participants [of the studies they were conducting] and thereby ‘to intensify the generation of alternative hypotheses.’” In this way, additional alternatives may

²³ See (Conee 2005: 66) and (Davis 2007: 436) for the idea of separating the account of our reactions to skeptical arguments and the account of bank case reactions. See (Sosa 1999: 148) and (Feldman 2001: 75f) for candidate fallacies that might be involved in skeptical arguments. Sosa, for example, argues that knowledge requires “safe” not “sensitive” beliefs. On that basis, he goes on to suggest that skeptical arguments may sound compelling because “[s]afety and sensitivity, being mutual contrapositives, are easily confused.”

²⁴ Sripada and Stanley (2012), Pinillos (2012) and Buckwalter and Schaffer (2013) claim to have confirmed such effects. The results from Pinillos, however, seem dubitable to me for the reasons mentioned in (Buckwalter and Schaffer 2013: Sec. 2.2) and (Buckwalter 2014: Sec. 3.2).

come to mind when one reads cases where the stakes are high. Thus, readers may be reminded that these alternatives are not unlikely. As a result, they may be less prone to ascribe knowledge because the epistemic position of the relevant subject will be assumed to be weakened (see the argument from reminding above.) Second, an explanation of the salience effect alone seems progress enough, so even if the above account of the stakes effect fails, this would not be a devastating result.

5. Conclusion

To sum up, projective invariantists can meet the challenge from the Contextualist Data. These data, then, provide no reason to reject invariantism.

Acknowledgements: Thanks to two anonymous referees, Catherine Diehl, Beate Krickel, Thomas Krödel, Dan López de Sa, Stefan Reining, Julia Zakkou, the participants of Geert Keil's and Tobias Rosefeldt's research seminars and audiences in Berlin, Dubrovnik and Konstanz for very helpful comments on earlier versions of the paper.

Bibliography

- Bach, K. (2005), 'The Emperor's New "Knows"', in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy. Knowledge, Meaning, and Truth* (Oxford: Oxford University Press), 51–90.
- (2008), 'Applying Pragmatics to Epistemology', *Philosophical Issues*, 18/1: 68–88.
- (2010), 'Knowledge In and Out of Context', in J. K. Campbell, M. O'Rourke, and H. Silverstein (eds.), *Knowledge and Skepticism* (Topics in Contemporary Philosophy, Cambridge, MA: MIT Press), 105–136.
- Baumann, P. (2011), 'WAMs: Why Worry?', *Philosophical Papers*, 40/2: 155–177.
- Blome-Tillmann, M. (2013), 'Knowledge and Implicatures', *Synthese*, 190/18: 4293–4319.
- Buckwalter, W. (2014), 'The Mystery of Stakes and Error in Ascriber Intuitions', in J. Beebe (ed.), *Advances in Experimental Epistemology* (Advances in Experimental Philosophy, New York, NY: Bloomsbury Publishing).
- Buckwalter, W., and Schaffer, J. (2013), 'Knowledge, Stakes, and Mistakes', *Noûs*, 2013: 1–34.
- Conee, E. (2005), 'Contextualism Contested Some More', in M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology* (Contemporary Debates in Philosophy, 3, Malden, MA: Blackwell Pub.), 62–66.
- Davis, W. A. (2007), 'Knowledge Claims and Context: Loose Use', *Philosophical Studies*, 132/3: 395–438.
- DeRose, K. (2009), *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. I* (Oxford: Oxford University Press).
- (2011), 'Contextualism, Contrastivism, and X-Phi Surveys', *Philosophical Studies*, 156/1: 81–110.
- Dougherty, T., and Rysiew, P. (2009), 'Fallibilism, Epistemic Possibility, and Concessive Knowledge Attributions', *Philosophy and Phenomenological Research*, 78/1: 123–132.
- Fantl, J., and McGrath, M. (2009), *Knowledge in an Uncertain World* (Oxford: Oxford University Press).

- Feldman, R. (2001), 'Skeptical Problems, Contextualist Solutions', *Philosophical Studies*, 103/1: 61–85.
- Gerken, M. (2011), 'Epistemic Focal Bias', *Australasian Journal of Philosophy*, 91/1: 41–61.
- (2012), 'On the Cognitive Basis of Knowledge Ascriptions', in J. Brown and M. Gerken (eds.), *Knowledge Ascriptions* (Oxford: Oxford University Press), 140–170.
- Hansen, N., and Chemla, E. (2013), 'Experimenting on Contextualism', *Mind & Language*, 28/3: 286–321.
- MacFarlane, J. (2005), 'The Assessment Sensitivity of Knowledge Attributions', in T. S. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology. Volume 1* (Oxford Studies in Epistemology, 1, Oxford: Oxford University Press), 197–234.
- Nagel, J. (2010), 'Knowledge Ascriptions and the Psychological Consequences of Thinking about Error', *The Philosophical Quarterly*, 60/239: 286–306.
- Nagel, J., Juan, V. S., and Mar, R. A. (2013), 'Lay denial of knowledge for justified true beliefs', *Cognition*, 129/3: 652–661.
- Pinillos, N. Á. (2012), 'Knowledge, Experiments and Practical Interests', in J. Brown and M. Gerken (eds.), *Knowledge Ascriptions* (Oxford: Oxford University Press).
- Rysiew, P. (2001), 'The Context-Sensitivity of Knowledge Attributions', *Noûs*, 35/4: 477–514.
- Schaffer, J. (2004), 'From Contextualism to Contrastivism', *Philosophical Studies*, 119/1/2: 73–103.
- Schaffer, J., and Knobe, J. (2012), 'Contrastive Knowledge Surveyed', *Noûs*, 46/4: 675–708.
- Sosa, E. (1999), 'How to Defeat Opposition to Moore', *Philosophical Perspectives*, 33/13: 141–153.
- Sripada, C. S., and Stanley, J. (2012), 'Empirical Tests of Interest-Relative Invariantism', *Episteme: A Journal of Social Epistemology*, 2012: 3–26.
- Stanley, J. (2005), *Knowledge and Practical Interests* (Short Philosophical Books, Oxford: Oxford University Press).

Williamson, T. (2005), 'Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge', *The Philosophical Quarterly*, 55/219: 213–235.