Philos Stud (2018) 175:2841–2857
https://doi.org/10.1007/s11098-017-0984-4

Anti-intellectualism, egocentrism and bank case intuitions

Alexander Dinges[1]

Published online: 27 September 2017

Abstract Salience-sensitivity is a form of anti-intellectualism that says the following: whether a true belief amounts to knowledge depends on which error-possibilities are salient to the believer. I will investigate whether salience-sensitivity can be motivated by appeal to bank case intuitions. I will suggest that so-called third-person bank cases threaten to sever the connection between bank case intuitions and salience-sensitivity. I will go on to argue that salience-sensitivists can overcome this worry if they appeal to egocentric bias, a general tendency to project our own mental states onto others. I will then suggest that a similar strategy is unavailable to stakes-sensitivists, who hold that whether a true belief amounts to knowledge depends on what is at stake for the believer. Bank case intuitions motivate salience- but not stakes-sensitivity.

Keywords Anti-intellectualism · Bank cases · Egocentric bias · Epistemic contextualism

1 Introduction

Intellectualism is the following view: whether a true belief amounts to knowledge depends only on "truth-relevant" factors, that is, on factors affecting "how likely it is that the belief is true, either from the point of view of the subject or from a more objective vantage point" (DeRose 2009: 24). Let's refer to the truth-relevant properties of a belief as the believer's *epistemic position* (regarding the belief in question). Intellectualism then is the view that whether a true belief amounts to knowledge depends only on the believer's epistemic position. *Anti-intellectualism*,

✉ Alexander Dinges
alexander.dinges@uni-hamburg.de

[1] Philosophisches Seminar, Universität Hamburg, Von-Melle-Park 6, 20146 Hamburg, Germany

on its most general construal, is the negation of intellectualism: whether a true belief amounts to knowledge depends not only on the believer's epistemic position. What I will call s*alience-sensitivity* is a specific form of anti-intellectualism. According to *salience-sensitivity*, whether a true belief amounts to knowledge depends on which error-possibilities are salient to the believer. Note that the dependence in play here is metaphysical, not causal (see below for some further discussion).

The most naïve way of implementing salience-sensitivity would be this: Knowledge is a matter of ruling out a restricted set of error-possibilities. Salient error-possibilities are the possibilities a subject attends to. S knows that p only if S can rule out all error-possibilities salient to S in this sense.

Such a view might make knowledge too fleeting. For instance, we would lose virtually all our knowledge by just thinking about the possibility of being a brain in a vat (assuming that we cannot rule out this possibility).[1] Note though that salience-sensitivists can introduce various kinds of constraints on when an error-possibility has to be ruled out to obtain knowledge even given its salience. For instance, they could say that the likelihood of the error-possibility must lie above a certain threshold. This threshold in turn could be fixed once and for all. Alternatively, it might vary with parameters such as what is at stake (depending on how we deal with issues such as the stakes versions of the bank cases to be described below). Salience-sensitivists could also introduce more demanding notions of salience. For instance, they could say that an error-possibility is salient only if one takes it seriously in some sense.[2]

One might worry that knowledge remains too fleeting even given such restrictions. This worry will have to be addressed at some point, but I want to set it aside for now. My question will be whether salience-sensitivity is motivated in the first place. If there is a good motivation for the view, we might be more forgiving regarding the indicated concern.

Salience-sensitivity is at best a minority position in the recent debate.[3] In what follows, I want to show that this status is undeserved. There are good reasons to adopt the view. One prominent way of arguing for salience-sensitivity starts from lottery considerations and the skeptical puzzles they give rise to. Salience-sensitivity appears to be at least one possible way of dealing with these puzzles.[4] Here I will be concerned with a different motivation. I want to show that salience-sensitivity can be motivated by appeal to intuitions about familiar scenarios such as the bank cases. Bank case intuitions are puzzling, and salience-sensitivity can contribute to an explanation of how they come about.

This motivation can be questioned on different ground. First, one may doubt the intuitions. There is an ongoing debate on this issue in the recent experimental philosophy literature. I won't go into the details here and take the intuitions for

[1] See e.g. Blome-Tillmann (2009: 246).

[2] See e.g. Lewis (1996) and Blome-Tillmann (2009) for further discussion.

[3] Hawthorne (2004) seems to be the only one to endorse it.

[4] See Hawthorne (2004).

granted instead. As I see it, they are supported by recent studies.[5] A second important problem remains. So-called third-person bank cases threaten to sever the connection between bank case intuitions and salience-sensitivity. This worry will be addressed in detail in this paper. In particular, I will argue that salience-sensitivists can overcome it if they appeal to egocentric bias, a general tendency to project our own mental states onto others.

The previous considerations naturally raise the question of whether a similar strategy is available for the much more prominent anti-intellectualist cousin of salience-sensitivity, *stakes-sensitivity*. According to *stakes-sensitivity*, whether a true belief amounts to knowledge depends on what is at stake for the believer.[6] I will address this issue and suggest that, even though stakes-sensitivists have appealed to egocentrism in the literature, the strategy is problematic in a way that doesn't affect egocentric salience-sensitivist accounts. As things stand, bank case intuitions can be used to motivate salience- but not stakes-sensitivity.

Note that the question of whether knowledge is sensitive to stakes or salient alternatives is of utmost philosophical importance. Anti-intellectualism has been variously employed to solve skeptical puzzles and to explain intuitive connections between knowledge and action.[7] Whether anti-intellectualism can be so employed, though, crucially depends on which non-epistemic factor is relevant for knowledge. Skeptical puzzles presumably require sensitivity to salient alternatives. Knowledge-action principles presumably require sensitivity to stakes. I show that only the former sensitivity is independently supported by bank case intuitions.

The plan for the paper is as follows: First, I will explain how third-person bank cases challenge salience-sensitivity based accounts of bank case intuitions. Second, I will show that salience-sensitivists can respond to this concern by appeal to egocentric bias. Finally, I will show that a similar strategy is unavailable to stakes-sensitivists.

2 Salience-sensitivity and third-person cases

Here is a pair of first-person bank cases salience-sensitivists might put forth in support of their view. The description begins with a common case-setup that eventually branches up into two cases (the labels are supposed to refer to the case-setup together with the respective branch).

> Hannah and Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday

[5] See below for references.

[6] See e.g. Stanley (2005) and Fantl and McGrath (2009) for proponents of such a view. Knowledge may be seen as depending on other practical factors beyond what is at stake (such as time constraints). See e.g. Shin (2014). I am focusing on stakes effects because they are most familiar. What I have to say about stakes should straightforwardly carry over to other putative practical determinants of knowledge.

[7] See Hawthorne (2004) for pertinent discussion of skeptical puzzles and e.g. Fantl and McGrath (2009) for knowledge-action principles.

afternoons. Hannah and Sarah consider their options. They could either queue or come back tomorrow, on Saturday, before they go to town for shopping. The question is whether the bank will be open tomorrow. Hannah has a clear recollection of depositing a paycheck at the bank 2 weeks before on a Saturday. She informs Sarah about this.

PLAIN Then she says, "I know the bank will be open tomorrow. Let's come back then."

ERROR Sarah points out that the bank may have changed its hours in the meantime. Hannah hasn't checked that. She says, "I don't know the bank will be open tomorrow. Let's queue."

Let's stipulate that Hannah truly believes that the bank will be open in both cases. Intuitively, Hannah's knowledge ascription in PLAIN as well as her knowledge denial in ERROR are correct.[8] But this is puzzling because Hannah's epistemic position seems to be the same in PLAIN and ERROR. In both cases she remembers having been at the bank two weeks before.[9]

Salience-sensitivity appears to solve this puzzle. Whether Hannah's true belief that the bank will be open amounts to knowledge depends not only on her epistemic position, but also on which error-possibilities are salient to her. She considers an additional error-possibility in ERROR. Hence we can unproblematically say that she knows the bank will be open in one case but not the other. This might seem like a strong motivation for salience-sensitivity: salience-sensitivity explains an otherwise puzzling pattern of intuitions.

Appealing as this reasoning may seem, third-person versions of the bank cases threaten to undermine it. Here are third-person versions of PLAIN and ERROR:

Hannah and Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Hannah and Sarah consider their options. They could either queue or come back tomorrow, on Saturday, before they go to town for shopping. The question is whether the bank will be open tomorrow. Hannah calls up Bill on her cell phone and asks Bill whether the bank will be open on Saturday. Bill replies by telling Hannah, "Well, I was there two weeks ago on a Saturday, and it was open." Hannah informs Sarah about this.

[8] As indicated, I don't want to address the empirical issue of whether these intuitions are real. Note, however, that at least similar intuitions have been confirmed in a range of recent studies. See e.g. Schaffer and Knobe (2012), Nagel et al. (2013) and Buckwalter and Schaffer (2015). Earlier studies failed to confirm the effects. See Schaffer and Knobe (2012: Sect. 2) for an overview. See the more recent studies for candidate accounts of this divergence.

[9] One way to respond to bank case intuitions is to deny that Hannah's epistemic position remains the same throughout the cases. See e.g. Pinillos (2011: 682f) and Dinges (2016). I will grant that Hannah's epistemic position remains the same for the purposes of this paper. My goal here is to investigate whether salience- and stakes-sensitivity can help to explain bank case intuitions. I don't want to show that there are no other candidate accounts.

$PLAIN_{3rd}$ Then she says, "Bill knows the bank will be open tomorrow. Let's come back then."

$ERROR_{3rd}$ Sarah points out that the bank may have changed its hours in the meantime. Hannah (correctly) thinks that Bill hasn't checked that. She says, "Bill doesn't know the bank will be open tomorrow. Let's queue."

Let's stipulate that Bill truly believes the bank will be open in both cases. As before, Hannah's knowledge ascription in $PLAIN_{3rd}$ and her knowledge denial in $ERROR_{3rd}$ are intuitively correct.[10] But this time, salience-sensitivity doesn't get a handle on these intuitions. For the subject of the knowledge ascription is Bill. And Bill doesn't consider an additional error-possibility in $ERROR_{3rd}$ (only Hannah and Sarah do).

We need an alternative explanation of third-person case intuitions. What could such an alternative explanation look like? There is a range of options. We could endorse epistemic contextualism or relativism, warranted assertability maneuvers or one of the various psychological accounts of bank case intuitions on offer in the literature.[11] All these accounts face independent problems, but we need not delve into the details here. The problem for salience-sensitivity is simply this: to the extent that epistemic contextualism, relativism, etc. are viable accounts of third-person case intuitions, they are equally viable accounts of first-person case intuitions.[12] So if the salience-sensitivist appeals to any such theory, she already has an account of all bank case intuitions. Salience-sensitivity itself becomes an idle wheel. This means that salience-sensitivists can no longer claim to contribute towards solving the puzzle from bank case intuitions. Hence, these intuitions can no longer be used to motivate the view. This is the challenge for salience-sensitivity from third-person cases.

I will offer a novel way of addressing this challenge below. Before that, though, it is worthwhile to briefly consider existing attempts to deal with third-person cases on behalf of salience-sensitivity (and anti-intellectualism more generally), and why

[10] Blome-Tillmann (2008: 31), for instance, endorses these intuitions with respect to a third-person case pair featuring a zebra and the possibility of it being a painted mule. Corresponding experimental studies are hard to come by. Intuitions about third-person cases have so far been tested only in Feltz and Zarpentine (2010). They didn't find relevant difference regarding our intuitions about these cases. However, they didn't find relevant differences regarding our intuitions about first-person cases either. These latter differences have been confirmed in subsequent studies with improved methodology (see footnote 8). It seems reasonable to expect similar results for methodologically improved third-person cases. But this will have to be investigated further.

[11] See e.g. DeRose (2009) for a defense of contextualism, MacFarlane (2005) for relativism, Brown (2006) for a warranted assertability maneuver and Nagel (2010a, b) and Gerken (2013) for psychological accounts.

[12] See e.g. DeRose (2009: Sect. 2.7) for how contextualists deal with third-person cases, MacFarlane (2005: Sect. 4.1) for how relativists do so and Brown (2006: 425–427) for a pertinent warranted assertability maneuver. Gerken (2013: Sect. 5.2) describes his view in a way that is entirely neutral between first- and third-person cases. Nagel's account will be discussed in more detail below. It will become obvious that the account as a whole also applies to first- and third-person cases. I will argue though that salience-sensitivists can isolate one particular aspect of Nagel's account (namely, the appeal to egocentric bias) to plausibly underwrite their position.

they are problematic. Hawthorne (2004: 164) appeals to the so-called "availability heuristic" and its distorting influence on "our estimation of the likelihood of an event".[13] The availability heuristic is a heuristic we use to judge the likelihood of events. We basically match the likelihood we ascribe to a given event with its cognitive availability. This heuristic generally leads to decent estimates, but it may lead us astray when events are available for reasons unrelated to likelihood. This, Hawthorne suggests, happens when we read cases like $\text{ERROR}_{3\text{rd}}$. The mentioning of the error-possibility in this case makes error artificially more available. Hence, we judge it as more likely than it actually is. This means that we judge Bill's epistemic position to be weaker than it actually is. As a result, we mistakenly judge that he doesn't know the bank will be open. In $\text{PLAIN}_{3\text{rd}}$, we aren't misled in a similar way, hence we correctly judge that Bill knows the bank will be open.

The availability account has recently been criticized on empirical grounds.[14] We can leave this criticism aside. Just note the following two basic points. First, the availability account is entirely independent of salience-sensitivity. As indicated, the supposed difference between the bank cases concerns the perceived *likelihood*, not the *salience*, of error-possibilities (the availability heuristic affects likelihood judgements not salience). This means that, given the account, we can explain our intuitions even if we think that salience of alternatives has no role to play in determining knowledge. Second, the account applies to first-person cases just as much as it applies to third-person cases. For the mentioning of the error-possibility in ERROR should raise the perceived likelihood of error in just the same way in which it does in $\text{ERROR}_{3\text{rd}}$.[15] Given these two points, we can dismiss the availability account. Just like the previously mentioned accounts, it leaves no work to do for salience-sensitivity.[16]

The following account is more promising in this regard. Stanley (2005) offers it on behalf of stakes-sensitivity, but it could equally be employed by salience-sensitivists. When we assess whether Hannah's knowledge claims in the third-person cases are true, we assess whether "Bill would know, were he in Hannah and Sarah's practical situation" (Stanley 2005: 102). Given that salience-sensitivity is true, this yields an account of why we agree with both of Hannah's claims. Bill arguably knows that the bank will be open. Even according to salience-sensitivity, he wouldn't lose this knowledge if he considered the error-possibilities salient to Hannah in $\text{PLAIN}_{3\text{rd}}$. After all, in this case, Hannah doesn't consider error-

[13] See e.g. Tversky and Kahneman (1973) for a general account of the availability heuristic and e.g. Slovic et al. (1982) for the kind of distorting influence Hawthorne has in mind.

[14] See Nagel (2010b).

[15] See e.g. Cohen (2004: 489), MacFarlane (2005: 214), Williamson (2005: 226), Stanley (2005: 101) and Schaffer (2006: 92f) for this latter observation.

[16] Note that Hawthorne presents the availability account only as an account of the appropriateness of knowledge *denials* in third-person cases. He offers a different account for the *in*appropriateness of knowledge *ascriptions*. This account does appeal to salience-sensitivity (and is further based on the knowledge norm of assertion and the factivity of knowledge). See Hawthorne (2004: 160). Again, we need not dwell on the viability of this account. The important point for our purposes is just that the availability account makes it redundant. For an increased assumed likelihood of error explains why we don't ascribe knowledge just as much as why we deny knowledge.

possibilities that he cannot rule out. He would lose his knowledge though if he considered the error-possibilities salient to Hannah in ERROR_{3rd}. For, he cannot rule out the possibility of changed opening hours.

This account crucially relies on salience-sensitivity, for if we reject salience-sensitivity, then the question of whether Bill knows doesn't depend on whether he considers the error-possibilities salient to Hannah in PLAIN_{3rd} or ERROR_{3rd}. But is the account plausible? As it stands, it certainly looks unmotivated. Why should we think that, in assessing Hannah's claims, we assess whether Bill *would know* that the bank will be open if he were in Hannah's respective context? Why don't we just assess whether he knows this? After all, Hannah is making claims about whether Bill knows, not about whether he *would know* in this or that situation. Stanley has some things to say here, but he hasn't convinced many. DeRose (2009: 235), for instance, worries that, while Stanley's account "is *possible*, he provides, and I can see, little reason to actually accept it."[17] I agree that Stanley doesn't provide the required motivation for his view, and I have nothing to add to DeRose's considerations in this regard, so I will not elaborate on this issue further.[18]

At this stage, salience-sensitivists could just bite the bullet and grant that bank case intuitions do not contribute to the motivation of their view. This may not leave the position unmotivated (see the remarks on lottery considerations above). But it is a substantial concession. After all, it is notoriously difficult to explain bank case intuitions, so it is a major asset of salience-sensitivity if it can explain these intuitions. In what follows, I will argue that salience-sensitivists need not concede as much. Egocentric bias straightforwardly complements their position to yield a promising account of first- and third-person cases. More specifically, it provides the missing motivation for Stanley's just described account by yielding an explanation for why we assess Bill as if he shared Hannah's concerns.

3 Salience-sensitivity and egocentric bias

Nagel (2010b) proposes an intriguing psychological account of the presented bank case intuitions that doesn't appeal to anti-intellectualism. Let us briefly look at her account to clarify the contrast between it and the closely related salience-sensitivist alternative I will present below.

Nagel's basic idea regarding *first-person* cases is this: As a matter of psychology, subjects concerned with additional error-possibilities tend to make a greater effort

[17] See Schaffer (2006: 93) and Wright (2011: 107) for the same assessment. See also DeRose (2012: 706ff).

[18] See Wright (2011) and Kim (2015) for further attempts to deal with third-person cases on behalf of anti-intellectualism. Wright (2011: 108) effectively suggests a warranted assertability maneuver, according to which the knowledge denial in ERROR_{3rd} seems correct because it should "be understood as claiming that Bill does not have sufficient information to settle the question of whether the bank's hours have changed." We have seen already why such accounts are problematic. Kim's (2015: 5) proposal saddles the anti-intellectualist with the view that "the truth of knows-whether assertions depends upon the context of utterance but the truth of knows-that assertions does not." It may be possible to make this result palatable, but at least on the face of it, it seems like a bitter pill to swallow.

before they form beliefs, in particular, they tend to try to rule out the error-possibilities in question. This makes Hannah in ERROR appear hasty. By stipulation, she forms the belief that the bank will be open, but the effort she makes is less than what we expect of a subject in her predicament. Presumably any method of belief formation (including the one that Hannah employs, whatever it is) becomes less reliable when hastily employed. Since reliability is a truth-relevant parameter, we can conclude that Hannah's epistemic position will be assumed to be weaker in ERROR than in PLAIN. Hence we are at liberty to say that she knows in one case but not the other. As Nagel (2010b: 303) condensedly puts it (considering a case pair featuring the error-possibility of a red light shining on a white table),

> People who are actively thinking about the influence of lighting conditions on colour judgements can still go ahead and make their colour judgements without checking the lighting, but would typically do so only under conditions of compromised or motivated belief formation. […] But these conditions—haste, distraction, wishful thinking—are the sort of conditions that tend to lower accuracy of judgement. When the accuracy of one's judgement appears to be compromised, one seems to be a mere believer, rather than a knower.

How does this account apply to *third-person* cases? Bill in ERROR_{3rd} doesn't consider additional error-possibilities. Hence, one may think, there is no reason to consider him hasty when it is stipulated that he believes the bank will be open. To respond to this worry, Nagel (2010b: 301f) appeals to egocentric bias, or egocentrism for short.[19] We will consider this bias in more detail below, but the basic idea is that readers project their own worry about the salient error-possibility mentioned in ERROR_{3rd} onto Bill and thereby treat him as if he worried about that possibilities as well. As Nagel (2010b: 303) puts it, "once concerns about the possibility of tricky lighting have been raised for me, I illegitimately evaluate John B. Doe's predicament as if he shared those concerns." Given that, the previous story about first-person cases applies: Bill will be considered hasty because he doesn't address the error-possibilities taken to be salient to him.

These are the rough outlines of Nagel's account of bank case intuitions. The account to be presented in what follows relates to this account in the following way: On the one hand, it rejects Nagel's account of first-person cases and substitutes it by the salience-sensitivist account of first-person cases described above. As such, it avoids a range of concerns that have recently been raised against Nagel's account.[20] On the other hand, it maintains the appeal to egocentrism in order to deal with third-person cases.

To see that there is a genuine contrast between the two views in question, it is important to note that Nagel's account of first-person cases does *not* entail salience-

[19] See also Nagel (2008: 292).

[20] See e.g Fantl and McGrath (2009: Sect. 3.1.3), Sripada and Stanley (2012: Sect. 5), Pynn (2014: 129f) and Shin (2014: Sect. 7). Nagel (2012) herself suggests a problem with her account: it predicts that subjects in ERROR should not only fail to have knowledge, their perceived levels of justification and confidence should also be lowered. According to Nagel, though, empirical results do not support this prediction. We will see below that this is in line with salience-sensitivity.

sensitivity. On Nagel's account, knowledge plausibly depends on which error-possibilities are salient to the subject. For on her view, salient error-possibilities can *cause* us to judge a subject to be less reliable and hence to deny knowledge. This doesn't mean that she endorses salience-sensitivity. Salience-sensitivists posit a *metaphysical*, not just a *causal*, dependence of knowledge on salient error-possibilities. According to salience-sensitivity, knowledge depends on salient error-possibility in just the same way in which it depends, say, on whether the subject believes the proposition in question.

So much for Nagel's account. Let's turn to how salience-sensitivists can appeal to egocentrism to extend their view to third-person cases. As a first step, we need to get clearer on what egocentrism amounts to.

Nagel appeals to what she calls "epistemic egocentrism." Royzman et al. (2003: 38) provide an unending list of studies confirming the effect. They describe epistemic egocentrism as a bias consisting in

> a difficulty in […] setting aside […] information (knowledge) that one knows to be unattainable to the other party, with a result that one's prediction of another's perspective becomes skewed toward one's own privileged viewpoint.

We will see in a minute that *epistemic* egocentrism, on this understanding, is useless as far as bank case intuitions are concerned. Note for this reason that egocentrism extends beyond the epistemic.[21] It is not only difficult to set aside one's own privileged knowledge. Similarly, "it might be easy to confuse how salient something is to oneself with how salient it is to others." (Gilovich et al. 2000: 212) For instance, participants in a study were asked to don t-shirts they found embarrassing. They briefly entered a room with other study participants. After leaving the room, they estimated the number of people who were able to tell what was on their t-shirt. The estimated numbers were much higher than the actual numbers presumably because participants projected their own concern with the embarrassing t-shirt onto the onlookers, who were, in fact, much less concerned. As Gilovich et al. (2000: 214) put it, "[p]articipants wearing a potentially embarrassing t-shirt allowed their own (quite understandable) focus on the shirt to distort their estimates of how much it would command the attention of others."[22]

Salience-sensitivists can appeal to this latter effect to explain third-person bank case intuitions as follows: Just like we project concerns about an embarrassing t-shirt onto those who see us, readers of the bank cases may project the salience of the possibility that the bank has changed its hours onto Bill. They themselves attend

[21] I am not saying that Nagel used "epistemic egocentrism" in the way defined by Royzman et al. She presumably had the more general notion in mind. But since the studies she mentions support epistemic egocentrism only as defined by Royzman et al., the issue should be rectified. To be sure, the main difference between the position suggested here and Nagel's own account is not supposed to lie in different ways of implementing egocentrism. As indicated before, the main difference is that Nagel's account doesn't entail salience-sensitivity, while my account crucially does.

[22] For more on the so-called "spotlight effect," see e.g. Kenny and DePaulo (1993), Gilovich and Savitsky (1999), Savitsky et al. (2001), Epley et al. (2002) and Gilovich et al. (2002).

to this possibility through reading ERROR_{3rd}. Hence, they tacitly assume that Bill attends to this possibility as well.[23] The same doesn't go for PLAIN_{3rd}, where no error-possibility is mentioned. Correspondingly, the salience-sensitivist is free to say that Bill's epistemic position suffices (or will be assumed to suffice) for knowledge given the alternatives (assumed) to be salient to him in PLAIN_{3rd} but not ERROR_{3rd}.

Note that what is projected here is not *knowledge* but salient alternatives. That is why *epistemic* egocentrism doesn't give us what we want.

Two clarificatory remarks: First and most importantly, salience-sensitivity is far from an idle wheel in the presented account. If knowledge doesn't depend on salient error-possibilities, it is unclear why it should matter whether we egocentrically project what is salient to us onto Bill. Given the projection, Bill's situation resembles Hannah's situation in the first-person cases. But this doesn't get us anywhere unless an account of first-person cases is in place already. (That is also why Nagel offers the previously outlined independent account of first-person cases.) Second, the account outlined may be seen as an error-theory. We are presumably mistaken to project our concerns onto Bill. This in itself isn't a problem, for the indicated error is rooted in a general psychological bias we have to grant anyway.

Two empirical observations to commend the view outlined: First, the account has it that our intuitions are guided by the salience of the error-possibility *to us as readers*. We as readers project what is salient to us onto the relevant subject. This suggests that bank case intuitions should not depend on whether the salient error-possibility is mentioned or only considered in thought by a subject in the scenario or just part of the case description. Recent experimental findings underwrite this prediction.[24] Second, the account predicts that, while the subject in the bank cases knows the bank will be open in one case but not the other, the level of justification for this proposition remains the same. According to salience-sensitivity, the subject in the error-possibility scenario fails to know the bank will be open because, given the salient error-possibility, the constant level of justification no longer suffices for knowledge. Again, recent empirical results suggest that justification levels remain constant in this way.[25]

Let me now address three potential worries with the outlined approach. First, what about bank case pairs differing in what is at stake rather than salient error-possibilities (see below for such cases)? We have the same intuitions about these cases as we have about the cases presented above (see below for discussion). But

[23] Merely attending to the possibility may not suffice for salience. As indicated, maybe readers must also take the possibility seriously. This wouldn't seem to be a problematic constraint. Readers who don't take the error-possibility seriously presumably don't share the intuitions about the cases. Note that the fact that readers (unlike Hannah and Bill) can rule out the possibility of changed opening hours doesn't entail that they don't take this possibility seriously; at least not on the notions of taking something seriously that the salience-sensitivist might want to employ. According to salience-sensitivity, the possibilities we take seriously are precisely the ones that we have to rule out. Hence, being able to rule out a possibility must be compatible with taking it seriously.

[24] See Alexander et al. (2014).

[25] See Nagel (2012). Nagel's studies show slight differences in justification levels, but they could be due to effects such as those described in Pinillos (2011: 682f) and Dinges (2016). Compare footnote 9.

salience-sensitivity cannot account for them. Some independent account will have to be provided, but this threatens to make salience-sensitivity superfluous again.

It doesn't. Accounts of third-person cases generally carry over to first-person cases. Accounts of stakes-cases generally don't carry over to salience-cases. Stakes-sensitivity, for instance, yields a candidate account of stakes cases even though, by itself at least, it has no bearing on salience cases. The same goes for the psychological account of stakes cases in Nagel (2008), so-called "interest contextualism" as defended e.g. in McKenna (2014) and the "warranted assertability maneuver" suggested in Gerken (2015: 231f).

Second, as readers of the bank cases, we know the bank will be open because this is stipulated. Correspondingly, we also know that it hasn't changed its hours. Given *epistemic* egocentrism, one might think, we should project this knowledge onto Bill and hence accept only Hannah's knowledge ascription.[26]

To respond, it is important to note that egocentrism is not an all-or-nothing affair. Of course, if we just plainly assume that Bill can rule out the possibility of changed opening hours, it shouldn't matter whether or not we assume that this possibility is salient to him. One loses knowledge only if an error-possibility becomes salient *that one cannot rule out*. Egocentrism, though, is subtler than that (at least in adults). It is a matter of being "*skewed* toward one's own privileged viewpoint." (Royzman et al. 2003: 38, my emphasis) We set aside our own perspective, just not as much as we should. Given epistemic egocentrism, people's judgements about *all* bank cases should be skewed towards the judgements they would make if they were to assume that the subject knows the bank will be open and can rule out the error-possibility in question. For this knowledge is available to readers of all cases. The salience of the error-possibility in $\text{ERROR}_{3\text{rd}}$ should still be expected to have an effect. People suppress their privileged knowledge to some extent and to that extent they should be drawn to deny knowledge once they project the salience of the error-possibility onto Bill (given salience-sensitivity).

Third, the effects of psychological biases can usually be surmounted in one way or another. For instance, one might consider modifying $\text{ERROR}_{3\text{rd}}$ by stating explicitly that Bill, unlike Hannah and Sarah, doesn't think about the possibility of changed opening hours. Shouldn't this suffice to undermine egocentrism? And wouldn't our intuitions about the case stay the same? Neither question has a clear answer. Pertinent studies on the bank cases don't exist (as far as I can see). Moreover, egocentrism proves extremely robust. It occurs even when "[p]eople know that others may see things differently than they do" (Gilovich et al. 2000: 212). With respect to *epistemic* egocentrism, Royzman et al. (2003: 47) similarly observe that "people are likely to succumb to [epistemic egocentrism] not only involuntarily […] but unwittingly, that is, while holding on to the belief that they succeeded in putting aside their privileged knowledge." Further empirical work needs to be done. We need a clear sense of how to surmount egocentrism and the intuitions about correspondingly modified bank cases. But I will leave this for

[26] See Nagel (2010a: footnote 18) for this worry (and a candidate response that differs from mine).

another occasion. The present worry shows that the outlined account is open to empirical refutation (as it should be), not that it has been refuted already.[27]

In sum, salience-sensitivity, if correct, contributes to an appealing account of the notoriously puzzling bank case intuitions. If, as some have argued, it also helps to solve skeptical puzzles, this would yield a strong case in favor of the view.

Of course, the latter issue would have to be investigated further. Relatedly, we would have to weigh these benefits against the potential costs indicated in the introduction (regarding the fleetingness of knowledge). For instance, the following awkward sounding sentence would seem to come out true if salience-sensitivity holds: "Hannah in ERROR doesn't know the bank will be open. She would know it though if she stopped worrying about changed opening hours."[28] These issues require separate treatment. For now, I will turn to stakes-sensitivity—roughly the view that knowledge depends on what is at stake for the putative knower.

Given the previous discussion, one naturally wonders whether stakes-sensitivists can also appeal to egocentrism to deal with third-person cases. I will argue that they cannot. At least, the story would be more complicated than in the case of salience-sensitivity.

4 Stakes-sensitivity

Let me begin by briefly summarizing how the problem from third-person cases affects stakes-sensitivist accounts of bank case intuitions.

At first, stakes-sensitivity may seem motivated given our intuitions about first-person case pairs like the following:

> Hannah and Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Hannah and Sarah consider their options. They could either queue or come back tomorrow, on Saturday, before they go to town for shopping. The question is whether the bank will be open tomorrow.
>
> LOW It is not especially important to deposit the paychecks by tomorrow, but if Hannah and Sarah come back tomorrow and the bank is closed, they will have to come back again on Monday before work, which would be annoying.

[27] See Nagel (2010b: 304f) and Alexander et al. (2014: Study 4 and Conclusion) for further discussion on how to surmount egocentrism in bank case assessments. Alexander et al. present data going somewhat against the egocentric story, Nagel mentions potentially more favorable results. All authors admit that the available data are inconclusive and, at the moment, I have nothing to add to this (except for maybe further incentives to carry out the relevant experiments).

[28] See e.g. Hawthorne (2004: 166f).

HIGH It is very important to deposit the paychecks by tomorrow because an important bill is coming due, so if Hannah and Sarah come back tomorrow and the bank is closed, they will be in a *very* bad situation.

Hannah has a clear recollection of depositing a paycheck at the bank two weeks before on a Saturday. She informs Sarah about this and says,

LOW "I know the bank will be open tomorrow. Let's come back then."

HIGH "I don't know the bank will be open tomorrow. Let's queue."

According to many authors, Hannah's knowledge claims in both these cases are intuitively correct.[29] This may be puzzling at first because the only difference between LOW and HIGH seems to be that the stakes are lower in LOW than in HIGH.

Stakes-sensitivity may seem to explain the puzzling intuitions. If knowledge depends on what is at stake for the subject, then Hannah may know the bank will be open in LOW but not in HIGH. For, there is more at stake for her in the latter case.

As before, third-person cases threaten to undermine this account. Here is a relevant case pair. The first paragraphs are omitted because they are exactly the same as in the first-person cases.

[Case setup as before.]

LOW_{3rd} [Description of low stakes as before.]
$HIGH_{3rd}$ [Description of high stakes as before.]

Hannah calls up Bill on her cell phone and asks Bill whether the bank will be open on Saturday. Bill replies by telling Hannah, "Well, I was there two weeks ago on a Saturday, and it was open." Hannah informs Sarah about this. Then she says,

LOW_{3rd} "Bill knows the bank will be open tomorrow. Let's come back then."
$HIGH_{3rd}$ "Bill doesn't know the bank will be open tomorrow. Let's queue."

Hannah's knowledge claims seem intuitively correct.[30] But stakes-sensitivity, taken by itself at least, doesn't get a handle on these intuitions, for the stakes for Bill remain constant.

The anti-intellectualist response strategies on offer in the literature have been discussed already. All plausible strategies make stakes-sensitivity redundant. At this point, stakes-sensitivists, just like salience-sensitivists, could bite the bullet and grant that their account cannot be motivated by appeal to bank case intuitions.

[29] Empirical support for these intuitions is a bit harder to find than for the corresponding salience cases, but see Sripada and Stanley (2012), Pinillos (2012), Pinillos and Simpson (2014) and Gao (2015: 102f). See Buckwalter (2014), Buckwalter and Schaffer (2015) and Turri (2017) for less favorable results.

[30] Once more, it must be noted that there are scarcely any empirical studies on the relevant third-person cases. Compare footnote 10.

Again, this may not leave the position unmotivated; for instance, stakes-sensitivity may be motivated by metaphysical principles about the relation between knowledge and action.[31] Still it should count as a substantial concession. What do we do? The appeal to egocentrism may seem to come in handy.[32] But the situation is not as straightforward as one might think.

Remember that the basic idea behind the *salience*-sensitivist egocentrism account was this: The mental states of the readers of PLAIN_{3rd} and ERROR_{3rd} differ. In particular, the possibility of changed opening hours is salient to them only when they read the latter scenario. This means that Bill's mental states are assumed to differ in the same way (by egocentrism). Hence, he will be assumed to know the bank will be open in one case but not the other (by salience-sensitivity).

Can the stakes-sensitivist adopt a similar account of the intuitions concerning LOW_{3rd} and HIGH_{3rd}? This seems problematic. The stakes-sensitivist will ultimately have to say that readers of the bank cases assume that more is at stake for Bill in HIGH_{3rd} than in LOW_{3rd}. Only this will yield a vantage point for stakes-sensitivity. This means that there must be a difference in what is at stake for the readers of these cases. For only then will egocentrism lead them to make similar assumptions about Bill. But this last premise is problematic for two reasons.

First, it is unmotivated. In particular, it is unclear why the stakes should rise for people reading a purely fictional story involving a high stakes subject.

One might suggest that we sympathize with Hannah in the bank cases just like we sympathize with other fictional characters. Given that it is very important for Hannah to be right, it might thus become important for us to be right.[33] This is unconvincing though. Hannah is so much underdescribed that it is hard to see how she could win our sympathies. Loewenstein and Small (2007: 120), for instance, make out a range of factors determining our sympathies for others. They claim that we sympathize with subjects "who share our own affective state, who are geographically or socially proximate, who are similar to us or are presented to us in a vivid fashion". It doesn't seem that any of this applies straightforwardly to Hannah. Maybe one could argue that readers assume by default that Hannah is supposed to be "similar" to them because nothing else is said about her. But further empirical research would be required to underwrite this hypothesis.

Another option would be that we somehow simulate Hannah's perspective when assessing her claims and then project her simulated concerns onto Bill via egocentrism.[34] Again, though, this account is at least underdeveloped. First, it is unclear why we should simulate Hannah rather than Bill. After all, to assess the truth-value of Hannah's claim, we'll have to assess Bill's not Hannah's state of mind. Second, given egocentrism, we project our concerns onto others. It is at best

[31] See e.g. Fantl and McGrath (2009) and Brown (2013: 244ff).

[32] Fantl and McGrath (2009: 56) mention this idea. As indicated, I think Stanley's (2005: 102f) account is most plausibly interpreted in terms of egocentrism too.

[33] Fantl and McGrath (2009: 56) and Nagel (2010a: 425) gesture towards accounts of the latter sort. See also Buckwalter and Schaffer (2015: 222), who mention Loewenstein and Small (2007) for pertinent psychological results.

[34] See Buckwalter and Schaffer (2015: 222) for a similar suggestion.

unclear, though, whether we also project merely simulated concerns, as the present account would require.

Even if the indicated premise can be underwritten after all, a second problem remains. If reading about fictional high stakes raises actual stakes, then knowledge becomes much more fleeting than the stakes-sensitivist presumably wants it to be. We would lose knowledge not only when we face an important (actual) decision, but already when presented with a short bank-case-like story explaining how pressing the issue in question is for some fictional character. For instance, we should get dialogues along the following lines:

A: I know the bank will be open.
B: Do you? Let me tell you a story about this person called "Hannah". She also thought she knew this, but it was *extremely* important for her to be right.
A: Ok, ok. I don't know it.

A's response seems pretty absurd.

Importantly, neither of these problems affects salience-sensitivity. First, whatever exactly salience amounts to, it seems clear that one way of making an error-possibility salient is by presenting it in the context of a short fictional story along the lines of the bank cases. This means, second, that we can lose knowledge by being offered such a story. But this prediction doesn't seem to be any worse at least than the prediction that we can lose knowledge by being directly told that the error in question could occur. Salience-sensitivists are bound to this latter prediction anyways.

The previous considerations may not be knock-down arguments against egocentric accounts of stakes cases. They should still suffice to show that, even when combined with egocentrism, stakes-sensitivity doesn't explain intuitions about third-person cases yet. We lack a plausible account of how the readers of the high stakes cases end up in high stakes contexts. This means that stakes-sensitivists still haven't offered an account of third-person bank case intuitions that doesn't make their view redundant.

5 Conclusion

Stakes-sensitivity gains no support from bank case intuitions. All available accounts of third-person cases make the view redundant. Salience-sensitivity does. When conjoined with egocentrism, it yields a promising account of first- and third-person cases. This is one point in favor of the view.

Acknowledgements I am grateful to Jie Gao, Dirk Kindermann, Jennifer Nagel, Jonathan Schaffer, the reading group Sprachphilosophie Berlin (most notably Emanuel Viebahn and Julia Zakkou), the members of the Forschungskolloquium in Hamburg, audiences in St Andrews and Cambridge and an anonymous referee for very helpful comments on earlier drafts of this paper.

References

Alexander, J., Gonnerman, C., & Waterman, J. (2014). Salience and epistemic egocentrism: An empirical study. In J. Beebe (Ed.), *Advances in experimental epistemology* (pp. 97–118). New York, NY: Bloomsbury Publishing.

Blome-Tillmann, M. (2008). The indexicality of "Knowledge". *Philosophical Studies, 138*(1), 29–53.

Blome-Tillmann, M. (2009). Knowledge and presuppositions. *Mind, 118*(470), 241–294.

Brown, J. (2006). Contextualism and warranted assertibility manoeuvres. *Philosophical Studies, 130*(3), 407–435.

Brown, J. (2013). Experimental philosophy, contextualism and SSI. *Philosophy and Phenomenological Research, 86*(2), 233–261. doi:10.1111/j.1933-1592.2010.00461.x.

Buckwalter, W. (2014). The mystery of stakes and error in ascriber intuitions. In J. Beebe (Ed.), *Advances in experimental epistemology* (pp. 145–173). New York, NY: Bloomsbury Publishing.

Buckwalter, W., & Schaffer, J. (2015). Knowledge, stakes, and mistakes. *Noûs, 49*(2), 201–234.

Cohen, S. (2004). Knowledge, assertion, and practical reasoning. *Philosophical Issues, 14*(1), 482–491.

DeRose, K. (2009). *The case for contextualism: Knowledge, skepticism, and context* (Vol. 1). Oxford: OUP.

DeRose, K. (2012). Replies to Nagel, Ludlow, and Fantl and McGrath. *Philosophy and Phenomenological Research, 84*(3), 703–721. doi:10.1111/j.1933-1592.2012.00586.x.

Dinges, A. (2016). Epistemic invariantism and contextualist intuitions. *Episteme, 13*(2), 219–232.

Epley, N., Savitsky, K., & Gilovich, T. (2002). Empathy neglect. Reconciling the spotlight effect and the correspondence bias. *Journal of Personality and Social Psychology, 83*(2), 300–312.

Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford: OUP.

Feltz, A., & Zarpentine, C. (2010). Do you know more when it matters less? *Philosophical Psychology, 23*(5), 683–706.

Gao, J. (2015). Advances in experimental epistemology. *International Studies in the Philosophy of Science*, *29*(1), 101–105.

Gerken, M. (2013). Epistemic focal bias. *Australasian Journal of Philosophy, 91*(1), 41–61.

Gerken, M. (2015). How to do things with knowledge ascriptions. *Philosophy and Phenomenological Research, 90*(1), 223–234.

Gilovich, T., Kruger, J., & Medvec, V. H. (2002). The spotlight effect revisited: Overestimating the manifest variability of our actions and appearance. *Journal of Experimental Social Psychology, 38*(1), 93–99.

Gilovich, T., Medvec, V. H., & Savitsky, K. (2000). The spotlight effect in social judgment: An egocentric bias in estimates of the salience of one's own actions and appearance. *Journal of Personality and Social Psychology, 78*(2), 211–222.

Gilovich, T., & Savitsky, K. (1999). The spotlight effect and the illusion of transparency: Egocentric assessments of how we are seen by others. *Current Directions in Psychological Science, 8*(6), 165–168.

Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: OUP.

Kenny, D. A., & DePaulo, B. M. (1993). Do people know how others view them? An empirical and theoretical account. *Psychological Bulletin, 114*(1), 145–161.

Kim, B. (2015). In defense of subject-sensitive invariantism. *Episteme*, FirstView, 1–19.

Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy, 74*(4), 549–567.

Loewenstein, G., & Small, D. A. (2007). The Scarecrow and the Tin Man: The vicissitudes of human sympathy and caring. *Review of General Psychology, 11*(2), 112–126.

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: OUP), pp. 197–234.

McKenna, R. (2014). Normative scorekeeping. *Synthese, 191*(3), 607–625.

Nagel, J. (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy, 86*(2), 279–294.

Nagel, J. (2010a). Epistemic anxiety and adaptive invariantism. *Philosophical Perspectives, 24*(1), 407–435.

Nagel, J. (2010b). Knowledge ascriptions and the psychological consequences of thinking about error. *Philosophical Quarterly, 60*(239), 286–306.

Nagel, J. (2012). Mindreading in Gettier cases and skeptical pressure cases. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions* (pp. 171–191). Oxford: OUP.
Nagel, J., Juan, V. S., & Mar, R. A. (2013). Lay denial of knowledge for justified true beliefs. *Cognition, 129*(3), 652–661.
Pinillos, N. Á. (2011). Some recent work in experimental epistemology. *Philosophy Compass, 6*(10), 675–688.
Pinillos, N. Á. (2012). Knowledge, experiments and practical interests. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions* (pp. 192–219). Oxford: OUP.
Pinillos, N. Á., & Simpson, S. (2014). Experimental evidence supporting anti-intellectualism about knowledge. In J. Beebe (Ed.), *Advances in experimental epistemology* (pp. 9–44). New York, NY: Bloomsbury Publishing.
Pynn, G. (2014). Unassertability and the appearance of ignorance. *Episteme, 11*(02), 125–143.
Royzman, E. B., Cassidy, K. W., & Baron, J. (2003). "I know, you know". Epistemic egocentrism in children and adults. *Review of General Psychology, 7*(1), 38–65.
Savitsky, K., Epley, N., & Gilovich, T. (2001). Do others judge us as harshly as we think? Overestimating the impact of our failures, shortcomings, and mishaps. *Journal of Personality and Social Psychology, 81*(1), 44–56.
Schaffer, J. (2006). The irrelevance of the subject: Against subject-sensitive invariantism. *Philosophical Studies, 127*(1), 87–107.
Schaffer, J., & Knobe, J. (2012). Contrastive knowledge surveyed. *Noûs, 46*(4), 675–708.
Shin, J. (2014). Time constraints and pragmatic encroachment on knowledge. *Episteme, 11*(2), 157–180.
Slovic, P., Fischhoff, B., & Lichtenstein, S. (1982). Fact versus fears. Understanding perceived risk. In D. Kahneman, P. Slovic, and A. Tversky (Eds.), *Judgment under uncertainty. heuristics and biases* (Cambridge: CUP), pp. 463–490.
Sripada, C. S., & Stanley, J. (2012). Empirical tests of interest-relative invariantism. *Episteme, 9*(1), 3–26.
Stanley, J. (2005). *Knowledge and practical interests*. Oxford: OUP.
Turri, J. (2017). Epistemic contextualism: An idle hypothesis. *Australasian Journal of Philosophy, 95*(1), 141–156.
Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology, 5,* 207–232.
Williamson, T. (2005). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *Philosophical Quarterly, 55*(219), 213–235.
Wright, S. (2011). Knowledge and social roles: A virtue approach. *Episteme, 8*(1), 99–111.