



# Knowledge and non-traditional factors: prospects for doxastic accounts

Alexander Dinges<sup>1</sup>

Received: 1 May 2019 / Accepted: 6 February 2020 / Published online: 20 February 2020  
© The Author(s) 2020

## Abstract

Knowledge ascriptions depend on so-called non-traditional factors. For instance, we become less inclined to ascribe knowledge when it's important to be right, or once we are reminded of possible sources of error. A number of potential explanations of this data have been proposed in the literature. They include revisionary semantic explanations based on epistemic contextualism and revisionary metaphysical explanations based on anti-intellectualism. Classical invariantists reject such revisionary proposals and hence face the challenge to provide an alternative account. The most prominent strategy here appeals to Gricean pragmatics. This paper focuses on a slightly less prominent strategy, which is based on the idea that non-traditional factors affect knowledge ascriptions because they affect what the putative knower believes. I will call this strategy *doxasticism*. The full potential of doxasticism is rarely appreciated in the literature and numerous unwarranted concerns have been raised. The goal of this paper is to present the strongest form of doxasticism and then to point out the genuine limitations of this position. I also sketch a closely related, more promising account.

**Keywords** Belief · Epistemic contextualism · Epistemic invariantism · Knowledge ascriptions · Salient alternative effects · Stakes effects

## 1 Introduction

Knowledge ascriptions depend on *non-traditional factors*, i.e. the kinds of factors that have taken center stage in discussions originating from intuitions about DeRose's (1992) bank cases and related case pairs. For instance, we become less inclined to

---

✉ Alexander Dinges  
alexander.dinges@fau.de

<sup>1</sup> Friedrich-Alexander-Universität Erlangen-Nürnberg, Institut für Philosophie, Bismarckstr. 1, 91054 Erlangen, Germany

ascribe knowledge when it's important to be right, or once we are reminded of possible sources of error. A number of potential explanations of this data have been proposed in the literature. They include revisionary semantic explanations based on epistemic contextualism and revisionary metaphysical explanations based on anti-intellectualism. Classical invariantists reject such revisionary proposals and hence face the challenge to provide an alternative account. The most prominent strategy here appeals to Gricean pragmatics (e.g. Rysiew 2001). This paper focuses on a slightly less prominent strategy, which is based on the idea that non-traditional factors affect knowledge ascriptions because they affect what the putative knower believes. I will call this strategy *doxasticism*. The full potential of doxasticism is rarely appreciated in the literature and numerous unwarranted concerns have been raised. The goal of this paper is to present the strongest form of doxasticism and then to point out the genuine limitations of this position. I will also outline a related, more promising account.

I begin by briefly describing the effects of non-traditional factors on knowledge ascriptions, focusing specifically on stakes effects and salient alternative effects (Sect. 2). After sketching contextualism and anti-intellectualism as the kinds of positions doxasticists reject (Sect. 3), I lay out doxastic accounts of both stakes effects and salient alternative effects (Sect. 4). Along the way, I deflect many familiar concerns. I point out residual worries with doxasticism and contrast this view with what I take to be a more promising account (Sect. 5). Then I conclude (Sect. 6).

## 2 Data

This section briefly reviews some basic data on how knowledge ascriptions are sensitive to non-traditional factors. This sensitivity comes in the form of case-based intuitions had by philosophers in the field. More recently, non-traditional factor effects have been studied in experimental philosophy. I will focus on the results from experimental philosophy, as they yield a more intersubjective basis for the subsequent discussion. Various types of non-traditional factors have been claimed to affect knowledge ascriptions. Most important among them are salient alternatives and stakes, and I will focus on these parameters in what follows.<sup>1</sup> Stakes effects and salient alternative effects have been confirmed in a range of different experimental set-ups. In this section, I will present only those studies that are particularly amenable to a doxastic account, leaving more problematic data for the discussion below. Notice that it is somewhat controversial whether the available studies actually confirm stakes effects and salient alternative effects. I will largely (though not entirely) ignore these controversies. My primary goal in this paper is to discuss accounts of non-traditional factor effects, while assuming that these effects exist.

Consider salient alternative effects. Buckwalter (2014) constructed versions of the familiar bank cases (DeRose 1992) that start out with a common case setup and branch three times. The first branching distinguishes between low and high stakes, the second

---

<sup>1</sup> Shin (2014) further studies time-constraint effects. For so-called contrast effects, see e.g. Schaffer and Knobe (2012).

between whether an error-possibility is mentioned, and the third between whether knowledge is ascribed or denied in the story. We end up with eight different vignettes.

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

Since they do not have an impending bill coming due, and have plenty of money in their accounts, it is not important that they deposit their paychecks by Saturday.

Since they have an impending bill coming due, and have very little money in their accounts, it is very important that they deposit their paychecks by Saturday.

Hannah says, “I was just at this bank two weeks ago on a Saturday morning, and it was open till noon. Let’s leave and deposit our paychecks tomorrow morning.”

[*Bank-Plain*] Sarah replies, “So the bank will be open tomorrow?”

Hannah says, “I was just here, I know that the bank will be open on Saturday.”

Hannah says, “Maybe you’re right, I don’t know that the bank will be open on Saturday.”

[*Bank-Error*] Sarah replies, “Well, businesses do change their hours sometimes. Just imagine how frustrating it would be driving here tomorrow and finding the door locked.”

Hannah says, “I was just here, I know that the bank will be open on Saturday.”

Hannah says, “Maybe you’re right, I don’t know that the bank will be open on Saturday.” (Buckwalter, 2014: 154)

These vignettes were presented to study participants. Each participant was assigned to one of the eight stories. Participants were asked to assume that “as it turns out, the bank really was open for business on Saturday” (154). Then they answered the question: “When Hannah said, ‘I (*know/don’t know*) that the bank will be open on Saturday,’ is what she said true or false?” (154). Answers were given on a five-point scale (1 = false; 3 = in-between; 5 = true). It turned out that participants’ truth-ratings of the knowledge ascription in *Bank-Error* was lower than in *Bank-Plain* independently of the variation in stakes. Similarly, their truth rating of the knowledge denial was higher in *Bank-Error* than in *Bank-Plain* independently of variations in stakes (157). This confirms salient alternative effects. People become less inclined to ascribe knowledge, and more inclined to deny knowledge, when an error-possibility is made salient to them.<sup>2</sup>

Let’s turn to stakes effects. In a study by Pinillos (2012), each participant read one of the following case descriptions:

[*Typo-Low*] Peter, a good college student has just finished writing a two-page paper for an English class. The paper is due tomorrow. Even though Peter is a pretty good speller, he has a dictionary with him that he can use to check and make sure there are no typos. But very little is at stake. The teacher is just asking

<sup>2</sup> For similar results, see e.g. Schaffer and Knobe (2012), Nagel et al. (2013), Alexander et al. (2014), and Gerken et al. (2020).

for a rough draft and it won't matter if there are a few typos. Nonetheless Peter would like to have no typos at all.

[*Typo-High*] John, a good college student has just finished writing a two-page paper for an English class. The paper is due tomorrow. Even though John is a pretty good speller, he has a dictionary with him that he can use to check and make sure there are no typos. There is a lot at stake. The teacher is a stickler and guarantees that no one will get an A for the paper if it has a typo. He demands perfection. John, however, finds himself in an unusual circumstance. He needs an A for this paper to get an A in the class. And he needs an A in the class to keep his scholarship. Without the scholarship, he can't stay in school. Leaving college would be devastating for John and his family who have sacrificed a lot to help John through school. So it turns out that it is extremely important for John that there are no typos in this paper. And he is well aware of this. (199)

Participants in each group were asked to fill in the blank in the following text: “How many times do you think Peter [John] has to proofread his paper before he knows that there are no typos? \_\_ times” (199). It turns out that “subjects tended to give higher answers in the high stakes condition *Typo-High* (median = 5) than in the low stakes condition *Typo-Low* (median = 2)” (200). This confirms stakes effects on knowledge ascriptions. People seem to require more evidence for knowledge when the stakes are high.<sup>3</sup>

Let me briefly respond to one important objection to the idea that such “evidence-seeking” studies confirm stakes effects on knowledge ascriptions. Buckwalter and Schaffer (2015: 214) suggest that we aren't seeing stakes effects *on knowledge ascriptions*. Instead, they hold, we are seeing stakes effects on the modal “has to” in the prompt, which supposedly receives a deontic interpretation (see also Rose et al. 2019: 240). The details of this idea don't matter because there is straightforward experimental evidence against it. I ran a follow-up study to the above evidence-seeking study using length-matched versions of *Typo-Low* and *Typo-High* from Buckwalter and Schaffer (2015: 208–209). I replaced the just-described prompt with a prompt asking participants to complete the sentence “Peter will know there are no typos in his paper once he has proofread it \_\_ times.” All putative deontic modality is removed (e.g. “has to” no longer features in the prompt), so the deontic account predicts that stakes effects should disappear. This isn't what I found. As before, more reads were required in *Typo-High* than in *Typo-Low*.<sup>4</sup>

In sum, there is good evidence for both stakes effects and salient alternative effects on knowledge ascriptions. We seem less inclined to ascribe knowledge when it's important to be right and once an error-possibility is pointed out to us. Much more

<sup>3</sup> For similar results, see e.g. Pinillos (2012), Pinillos and Simpson (2014), Buckwalter and Schaffer (2015), Francis et al. (2019), and Dinges and Zakkou forthcoming.

<sup>4</sup> 101 participants were recruited through Prolific Academic (68% female, mean age 32). They were randomly assigned to one of the stories and then received the prompt described. 22 participants failed a simple attention check (recalling the length of the paper) and were thus excluded from the analysis. A Mann–Whitney test for the remaining 79 participants indicated that the number of required reads was greater for *Typo-High* (*Mdn* = 3) than for *Typo-Low* (*Mdn* = 2):  $U = 508.500, p = .007, r = -.30$  (medium effect).

could be said here. But in what follows, I will take this data for granted and turn to potential explanations instead.

### 3 Anti-intellectualism and contextualism

This section sketches explanations of non-traditional factor effects based on contextualism and anti-intellectualism. I discuss these views just to set them aside. They are the kinds of views doxasticists reject.<sup>5</sup>

Contextualism is, very roughly, the position that the term “know” expresses differently demanding epistemic relations depending on non-traditional factors in the context in which it is used. For instance, we can think of knowledge in terms of one’s ability to rule out error-possibilities. According to contextualism, the context of use determines just how many error-possibilities one needs to rule out to count as a “knower.” As applied to e.g. salient alternative effects, the idea would be that people become less inclined to ascribe “knowledge” once an error-possibility has been made salient to them because in such contexts, “know” comes to express a more demanding epistemic relation, one that requires the ability to rule out the salient error-possibility (see e.g. Wright 2017 for a recent overview of the discussion surrounding contextualism).

Anti-intellectualism is, very roughly, the position that the strength of the evidence entailed by knowledge shifts with non-traditional factors in the context of the putative knower (rather than the context of the knowledge ascription). Take the protagonist of the typo stories above. In *Typo-Low*, the stakes are low for him. According to anti-intellectualism, this means that knowledge entails only moderately strong evidence (e.g. two rounds of proofreading). In *Typo-High*, the stakes are high for him. Thus, knowledge entails stronger evidence (e.g. five rounds of proofreading). Responses to the cases supposedly reflect these variable evidential demands (see e.g. Weatherson 2017 for a recent overview of the discussion surrounding anti-intellectualism).

Notice that I have tied contextualism to salient alternative effects and anti-intellectualism to stakes effects. This pairing is typical but not mandatory. A contextualist can hold that what is at stake in a given context determines the relation expressed by “know” (e.g. Blome-Tillmann 2014: 14–15), while an anti-intellectualist can maintain that the evidential demands for knowledge shift as more error-possibilities become salient to the putative knower (e.g. Hawthorne 2004: 159; Dinges 2018a). Doxasticists on my understanding reject all of these positions.

### 4 Doxasticism

In this section, I present doxasticism in more detail. I begin by presenting what I take to be the basic principles in a doxastic account. I then lay out how these principles can be used to explain non-traditional factor effects in the studies reported above as well

<sup>5</sup> Notice that doxasticism need not be understood as entailing the denial of contextualism and anti-intellectualism. Indeed, some forms of doxasticism naturally entail anti-intellectualism (e.g. Weatherson 2005; Ganson 2008; Ross and Schroeder 2014). I won’t discuss these forms of doxasticism here.

as in other, more challenging contexts. Notice that one should distinguish doxasticism as applied to salient alternative effects from doxasticism as applied to stakes effects. One can endorse one of these accounts while rejecting the other. I will present them in tandem to make parallels visible. I will duly tear them apart later when discussing objections.<sup>6</sup>

#### 4.1 Basic principles

The following principles are at the center of gravity in a doxastic account.

**SALIENCE** Unless we are rushed, distracted or otherwise impaired, we normally don't form a belief that *p* when an error-possibility to *p* is salient to us that we cannot rule out.

**STAKES** Unless we are rushed, distracted or otherwise impaired, we normally don't form a belief that *p* when we aren't in a position to act on the assumption that *p*.

SALIENCE, for instance, suggests that the protagonists in the bank cases will lose their belief that the bank will be open once the possibility of changed opening hours becomes salient to them. After all, they cannot rule it out. STAKES suggests that the protagonist of the typo stories will read his paper more frequently before he forms the belief that no typos remain when the stakes are high rather than low. This is because, in the high stakes context, he will need more evidence before he can act on the assumption that no typos remain i.e. before he can submit his paper. (The “unless”-clauses in SALIENCE and STAKES will become relevant later.)

Doxasticists can appeal to empirical data from psychology to justify *that* SALIENCE and STAKES hold. I will not rehearse this data here and just grant that it is sound.<sup>7</sup> Instead, I want to briefly address the question of *why* SALIENCE and STAKES hold. Doxasticists aren't committed to any specific view here. To clarify their position, though, it will be useful to outline their options.

<sup>6</sup> The doxastic account of stakes effects I will present most closely resembles (and may be identical to) the account suggested by Nagel (2008, 2010a). The doxastic account of salient alternative effects most closely resembles (and may be identical to) the account suggested by Nagel (2010b). For related accounts of stakes effects and salient alternative effects, see Bach (2005, 2008, 2010). For a precursor of the doxastic account of salient alternative effects, see Hawthorne (2004: 169).

<sup>7</sup> For a recent overview of the empirical evidence for STAKES, see Gao 2019: §2. See also Nagel (2008: §1), Nagel (2010a: §§2–3) and Nagel and Smith (2017: §3). The evidence available for SALIENCE is more elusive. Nagel (2010b: 303) states a principle along the lines of SALIENCE, but the empirical evidence she cites (Mayseless and Kruglanski 1987 and Kunda 1990) seems to bear primarily on STAKES. See Hawthorne (2004: 169), Nagel and Smith (2017: 100), and Dimmock (2019: 4) for comparatively unsupported statements of SALIENCE. Dimmock laudably mentions (Kelley 1972) as one possible source of evidence. Stronger evidence may be available. Dinges (2018b) cites ample empirical evidence suggesting that people lose confidence in a given hypothesis when specific sources of error are pointed out to them (see also Vogel 1990: 19–20; Hawthorne 2004: 162–166; and Williamson 2005: 226). This may support SALIENCE to the extent that, as seems plausible, people lose beliefs when they lose confidence. Notice that it may be questionable whether the available data supports precisely the principles SALIENCE and STAKES. As far as I can see, though, doxasticists need these principles to get their account going. So if they aren't supported by the data, then so much the worse for doxasticism.

One candidate explanation is normative. Doxasticists can hold that the conditions for *justified* belief become more demanding once the stakes rise or an error-possibility is made salient. For instance, they could adopt a principle along the following lines. S's belief that *p* is justified only if S's evidence for *p* is such that (i) it's good enough for S to act on *p* in her current context<sup>8</sup> and (ii) it rules out all error-possibilities to *p* currently salient to S. This would explain why we lose beliefs as the stakes rise and new error-possibilities become salient. The pertinent beliefs would simply cease to be justified. For instance, the protagonists in the bank cases may initially be justified to believe that the bank will be open but cease to be so justified once the possibility of changed opening hours is raised to salience. Similarly, the protagonist of the typo stories may be justified to believe that no typos remain after two rounds of proofreading when the stakes are low, whereas, when the stakes are high, he fails to be so justified because he cannot yet act on this assumption and thus submit his paper.

A caveat should be noticed. It's exceedingly natural to think that knowledge entails justified belief. Given the just-stated justification principle, knowledge will thus entail different evidential demands depending on non-traditional factors. When the stakes are high, for instance, knowledge will entail stronger evidence than when the stakes are low. This is, of course, a paradigmatic form of anti-intellectualism, a view that doxasticists (on my understanding) reject. So if the doxasticist adopts the above justification principle, she has to sever the link between knowledge and justified belief. For instance, she could argue that the evidential requirements for justified belief sometimes exceed the evidential requirements for knowledge.<sup>9</sup>

Another way to explain STAKES and SALIENCE is error-theoretic. Doxasticists can reject the variable justification principle above and hold instead that the evidential requirements for justified belief remain stable independently of non-traditional factors. They can thus retain the link between knowledge and justified belief. It also follows that our sensitivity to stakes and salient alternatives is somehow erroneous. Take the sensitivity to stakes to illustrate. Other things being equal, beliefs in low-stakes contexts and corresponding beliefs in high-stakes contexts will either both be justified, or they will both be unjustified. In the former case, our practice of giving up beliefs in the face of high stakes (following STAKES) would be mistaken. After all, it is a mistake to give up perfectly justified beliefs. In the latter case, our practice of maintaining beliefs in low stakes contexts would be mistaken because it is a mistake to hold on to unjustified beliefs. Doxasticists can endorse one of these error-theoretic commitments, and then provide some explanation of why we make the respective mistake in order to explain why STAKES and SALIENCE hold.

A mixed strategy is possible too. One might try to vindicate STAKES normatively, based on the idea that beliefs are justified only if the subject can act on the believed proposition, thus severing the link between knowledge and justified belief. Meanwhile, one can suggest that SALIENCE reflects erroneous belief formation processes (or vice versa).

<sup>8</sup> See e.g. Weatherston (2005), Ganson (2008), Fantl and McGrath (2009), and Ross and Schroeder (2014) for various ways to spell out and motivate conditions along the lines of (i).

<sup>9</sup> This may be plausible if we link belief to action and grant that the evidential requirements for action sometimes exceed the evidential requirements for knowledge (e.g. Brown 2008: §6).

All of the above accounts face challenges. It seems troublesome to sever the link between knowledge and justified belief; and error-theoretic commitments impose a burden of proof. Neither of these challenges, however, seems insurmountable at this point. Thus, I submit the above considerations as clarifications of the commitments of doxasticism, and as invitations to say more. But I won't treat them as an objection.

## 4.2 The basic account

Let's see how STAKES and SALIENCE can be used to explain non-traditional factor effects on knowledge ascriptions. Consider salient alternative effects first, and recall Buckwalter's factorial version of the bank cases. In *Bank-Plain*, the possibility of changed opening hours isn't salient. Thus, SALIENCE doesn't stand in the way of the protagonists forming the belief that the bank will be open, and doxasticists will hold that this belief amounts to knowledge. In *Bank-Error*, the possibility of changed opening hours is salient. Moreover, the protagonists can't rule it out. Thus, by SALIENCE, they won't normally form the belief that the bank will be open. Study participants assume that things are supposed to be normal in the story they read,<sup>10</sup> so they will assume that the protagonists don't form this belief. Since knowledge entails belief,<sup>11</sup> the protagonists will lack knowledge. According to the doxastic account, this explains why participants become less willing to ascribe knowledge.

Consider stakes effects, and recall Pinillos' typo stories. In *Typo-Low*, the stakes are low and the protagonist can arguably act on the assumption that the paper is free of typos after two rounds of proofreading. Thus, after two rounds, STAKES doesn't stand in the way of belief formation anymore, and the protagonist gains knowledge that no typos remain according to doxasticism. Meanwhile, in *Typo-High*, the protagonist should arguably read his paper five times before he can act on the assumption that no typos remain. Thus, five reads will normally be required for him to form the pertinent belief, given STAKES. Again, participants assume that things are supposed to be normal, and thus they assume that the protagonist will form the respective belief only after five rounds. Since knowledge entails belief, they say that he will know that the paper is free of typos only after five rounds.

Let me flag an implicit assumption in the outlined account. Even if STAKES and SALIENCE hold and even if knowledge entails belief, it doesn't follow that judgements about the typo cases and the bank cases differ between conditions. These principles merely suggest that, on a normal construal of these cases, there will be differences in knowledge. Participants could still fail to track these differences in their judgements. To bridge this gap, we have to assume that people track the joint predictions of STAKES and SALIENCE, on the one hand, and the belief requirement on knowledge, on the other, when assessing fictitious cases like the bank cases and the

<sup>10</sup> See Dafoe et al. (2018) for the more specific idea that participants fill in cases via Bayesian updating, and empirical support for this assumption.

<sup>11</sup> For (rare) opposition to this principle, see e.g. Radford (1966), Myers-Schulz and Schwitzgebel (2013), Farkas (2015), and Tebben (2019). Tebben suggests that knowledge entails "commitment" rather than belief. This would only strengthen the case for doxasticism because especially STAKES only becomes more plausible when applied to commitment rather than belief.



typo cases. I will challenge this *tracking assumption* later. For now, I just want to put it on the table.<sup>12</sup>

One immediate prediction of the doxastic account is nicely confirmed. On the doxastic account, we would expect that not only judgements about knowledge but also judgements about belief are affected by factors such as stakes and salient alternatives. After all, knowledge judgements supposedly vary because belief judgements do. Nagel et al. (2013: 656) correspondingly show that the salience of error-possibilities affects belief ascriptions. Similarly, we find stakes effects when “knows” is replaced by “believes” in the prompt from the typo studies above (Buckwalter 2014: 160–162; Buckwalter and Schaffer 2015: 209–214).<sup>13</sup> This is good news for the doxasticist, but the point shouldn’t be overstated. Anti-intellectualists who subscribe to the variable justification principle from above, for instance, make similar predictions, and thus the presented data equally confirms their position.

Some clarifications of the doxastic account are in order. First, doxasticism is distinct from anti-intellectualism. According to doxasticism, knowledge depends on non-traditional factors. In particular, non-traditional factors determine how much evidence is normally required to obtain knowledge by determining how much evidence is normally required for belief. Anti-intellectualism doesn’t follow because it doesn’t follow that knowledge *entails* different amounts of evidence depending on non-traditional factors. For instance, it is perfectly possible according to doxasticism to be in a high stakes situation, to have only moderate amounts of evidence and still to have knowledge. Such situations would merely be abnormal because we wouldn’t normally satisfy the belief condition for knowledge when we have only moderate amounts of evidence while the stakes are high.<sup>14</sup>

Second, doxasticism is not an error-theory. On this view, the protagonist of the bank cases, for instance, actually loses knowledge once an error-possibility has been raised to salience because she loses the underlying belief (on a normal construal of the cases). Similarly, the protagonist of the typo stories actually needs more evidence for knowledge when the stakes are high because he needs more evidence for belief (on a normal construal of the cases). We’ll see later in our discussion of “egocentric bias” (Sect. 4.5) that certain types of cases force some error-theoretic commitments on the doxasticist, but so far doxasticism accommodates our intuitions.

<sup>12</sup> Nagel (2008: 284, 2010a: 413) notices that such an assumption is required and provides empirical evidence for the assumption that our judgements about other people’s doxastic states track the predictions of STAKES and SALIENCE.

<sup>13</sup> These latter effects remain even when “knows” is replaced by “hopes” and “guesses” (Buckwalter 2014: 160–162; Buckwalter and Schaffer 2015: 209–214). On the doxastic account, this isn’t particularly surprising. Both “hopes” and “guesses” have one meaning at least which seems very similar to the meaning of “belief.” For instance, the point in time at which John in the typo cases would naturally assert things like “I believe that I’m done here” would also be the point in time at which he would naturally assert “I hope/guess I’m done here.” See also Nagel and Smith 2017: 98–99 on the idea that “epistemic intuitions can be elicited only by imagining the target agent as acting or asserting something.”

<sup>14</sup> Another way to put the difference between anti-intellectualism and doxasticism is as follows. According to anti-intellectualism, non-traditional factors partially constitute knowledge, while according to doxasticism, there are merely counterfactual or causal dependencies. See Ichikawa and Steup (2014: §12) for related thoughts.

Third, doxasticism is primarily concerned with the doxastic states of the protagonists of the cases discussed, not the doxastic states of the study participants. In maybe more familiar terms, doxasticists focus on the “subject” rather than the “ascriber” or the “assessor.” According to doxasticism, the doxastic states of the protagonists vary in line with STAKES and SALIENCE, and therefore their knowledge varies. Meanwhile, study participants simply track these facts. We’ll see later in our discussion of “egocentric bias” (Sect. 4.5) that specific cases force the doxasticist to focus on the study participant after all.

### 4.3 Stipulating belief

Doxasticism nicely accommodates the data presented so far, but it may appear insufficiently general. Schaffer and Knobe (2012: 695), for instance, confirm salient alternative effects, while explicitly stipulating that the protagonist’s confidence remains unaltered across conditions. Pinillos (2012: 203) similarly shows that stakes effects in his typo experiment remain when it is stipulated that the protagonist of the story forms the belief that there are no typos in his paper right after he finishes writing it. The doxastic account as outlined no longer gains traction. After all, it was based on the idea that shifts in knowledge judgements result from assumed shifts in belief. And it seems implausible that participants should make different assumptions about what the protagonists believe when the case descriptions explicitly hold this factor fixed.

An initially tempting way for doxasticists to respond is to say that the stipulated belief is unjustified when the stakes are high or an error-possibility is salient. This assumption could be based on the justification principle above, according to which high stakes and salient alternatives shift the evidential requirements for justified belief. Doxasticists could go on to claim that participants no longer want to ascribe knowledge because knowledge entails justified belief. We’ve seen though that this combination of views is unavailable to doxasticists. If knowledge entails justification and the evidential requirements for justification vary with non-traditional factors, anti-intellectualism follows. Our doxasticist, however, rejects anti-intellectualism.

Other strategies are available.<sup>15</sup> Recall that STAKES and SALIENCE contain an “unless”-clause. STAKES, for instance, says that we normally don’t form beliefs when we can’t act on them *unless* we are rushed, distracted or otherwise impaired. Now consider cases where it is stipulated that the protagonist forms the respective belief when the stakes are high or an error-possibility is salient. By STAKES and SALIENCE, this belief wouldn’t normally be formed *unless* the protagonist was rushed, distracted or otherwise impaired. Thus, participants will assume that the protagonist was rushed, etc. because they assume that the cases are supposed to be normal. This will make them less inclined to ascribe knowledge even if the belief condition is satisfied. After all, knowledge presumably entails that the belief was formed in a somewhat reliable way. And belief formation processes tend to become less reliable when employed while being rushed, etc. Take John in *Typo-High*, for instance. If he’s e.g. rushed, it becomes more likely that he misses typos when he proofreads his paper. Thus, he will

<sup>15</sup> The subsequent proposal is taken from Nagel (2008: 291), (2010a: 419), (2012: 680–681) (stakes effects) and (2010b: 303) and Nagel and Smith (2017: 100) (salient alternative effects).

be less reliable and less likely to obtain knowledge than his low stakes counterpart. Take the bank cases. A protagonist remembers her visit at the bank two weeks before. Our memory can be misleading though and when we are e.g. rushed, we are less likely to notice corresponding defeaters (such as that there was a sign on the door indicating that it was a special opening). This makes the respective belief formation process less reliable and thus less likely to yield knowledge.

A potential worry would be that we could simply stipulate that the protagonists in *Bank-Error/Typo-High* form their beliefs “no more hastily and with no more bias” than the protagonists in *Bank-Plain/Typo-Low* (Fantl and McGrath 2009: 45). One might suggest that, intuitively, this would leave stakes effects and salient alternative effects unaltered, contrary to what the indicated account predicts. These intuitions are unstable though. Nagel (2008: 293), for instance, maintains that “the inclination to ascribe knowledge to the low-stakes subject but not the high-stakes one does not persist when their cognitive situations are explained in full detail.” Further empirical research may help to determine whose intuitions are right, but I will not pursue this issue here.

Let me add that the described approach to stipulated-belief cases isn’t the only option for doxasticists. Stipulating belief may be ineffective for much more general reasons. Nagel and Smith (2017: 94), for instance, note “the limits of our powers of stipulation.” If it makes no sense in the story that the relevant protagonist forms the relevant belief (e.g. because the available evidence is way too weak), we may “instinctively” override this stipulation.<sup>16</sup> Relatedly, the term “belief” may be ambiguous and allow for weaker or stronger readings. Participants may choose a weaker reading to make sense of the fact that the protagonist forms a “belief” with otherwise insufficient evidence, while only a stronger reading is relevant for knowledge and governed by principles like STAKES and SALIENCE.<sup>17</sup>

#### 4.4 Being in a position to know

Pynn (2014: 130) suggests that non-traditional factor effects persist if we ask participants when the protagonist of the relevant story is *in a position to know* rather than when she knows. Doxasticism supposedly doesn’t explain this because belief is not required in order to be in a position to know. Even if a protagonist loses her belief, she may still be in a position to know.

To strengthen this worry, I confirmed the relevant intuitions in a follow-up study with the length-matched versions of *Typo-Low* and *Typo-High* from Buckwalter and Schaffer (2015: 208–209). I replaced “knows” by “is in a position to know” in the prompt from the previously reported study, where all deontic modality was removed. In line with the indicated intuitions, a stakes effect on the number of required reads

<sup>16</sup> This proposal has limitations. Nagel et al. (2013: 657) find salient alternative effects after all participants who didn’t ascribe the relevant belief have been filtered out. In this way, assumptions about belief were held fixed independently of a possibly overridable stipulation.

<sup>17</sup> See Pynn (2014: 129n10) for the latter observation. To accommodate the just-mentioned results from Nagel et al. (2013: 657), one would have to argue that “belief” denotes the weaker, irrelevant notion by default (see Hawthorne et al. 2016 for relevant discussion).

remained.<sup>18</sup> Notice further that “is in a position to know” isn’t just a philosophical term of art. A web search reveals that it is frequently used in ordinary English too. Of course, there may be differences between the ordinary and the technical usage. All that matters at this point, though, is that neither the technical nor the ordinary notion entails belief. And this assumption should be granted by everybody.

Doxasticists can respond to this concern as follows. The present objection implicitly assumes an analysis of the notion of being in a position to know along the following lines. S is in a position to know that p iff S’s overall situation is such that if she were to form the belief that p based on her evidence, then she would know that p. On this analysis, belief plays no relevant role. But it’s not clear why we should adopt this analysis in so far as we are trying to capture a folk concept. The following analysis seems at least equally promising. S is in a position to know that p iff S’s overall situation is such that she could easily come to know that p. This would entail that S could easily come to believe that p. And, given STAKES and SALIENCE, this latter condition will be harder to meet in high stakes cases or when an additional error-possibility has been made salient. Doxasticists can argue that non-traditional factor effects remain for this reason.

#### 4.5 Egocentric extensions

Another challenge for doxasticists comes from so-called “ignorant stakes cases” and cases where the error-possibility in question is made salient to the study participants but not the protagonists of the story. I will address this challenge in what follows.

Consider ignorant stakes cases. These are case pairs where the stakes vary as in e.g. *Typo-Low* and *Typo-High*. Additionally, it is stipulated that the protagonist is unaware of what is at stake. To illustrate, consider the following ignorant high stakes case from Pinillos (2012: 216–217):

[*Ignorant-Typo-High*] John, a good college student, has just finished writing a two-page paper for an English class. The paper is due tomorrow. Even though John is a pretty good speller, he has a dictionary with him that he can use to check and make sure there are no typos. There is a lot at stake. The teacher is a stickler and guarantees that no one will get an A for the paper if it has a typo. He demands perfection. John, however, finds himself in an unusual circumstance. He needs an A for this paper to get an A in the class. And he needs an A in the class to keep his scholarship. Without the scholarship, he can’t stay in school. Leaving college would be devastating for John and his family who have sacrificed a lot to help John through school. So it turns out that it is extremely important for John that there are no typos in his paper. However, John is unaware of what is really at stake. He thinks the teacher does not care at all if there are some or even many

<sup>18</sup> 100 participants were recruited through Prolific Academic (mean age 33, 62% female). They were randomly assigned to one of the stories and then received the prompt described. 23 participants failed a simple attention check (recalling the length of the paper) and were thus excluded in the analysis. A Mann–Whitney test for the remaining 77 participants indicated that participants required more reads in *Typo-High* ( $Mdn = 3$ ) than in *Typo-Low* ( $Mdn = 2$ ):  $U = 354.000$ ,  $p < .001$ ,  $r = -.47$  (large effect).

typos in the paper. Although John would like to have no typos, he is unaware that it would be extremely bad for him if there is but a single typo in the paper.

Such ignorant stakes cases also trigger stakes effects (e.g. Pinillos 2012: 202–203; Pinillos and Simpson 2014). On the doxastic account, this is puzzling. After all, protagonists can't adjust their doxastic states to their practical situation when they don't know about it.

An analogous worry can be raised for the doxastic account of salient alternative effects. Consider, for instance, the following case pair from Alexander et al. (2014: 99).

[*Furniture-Plain*] John A. Doe is in a furniture store. He is looking at a bright red table under normal lighting conditions. He believes the table is red.

[*Furniture-Error*] John B. Doe is in a furniture store. He is looking at a bright red table under normal lighting conditions. He believes the table is red. However, a white table under red lighting conditions would look *exactly* the same to him, and he has not checked whether the lighting is normal, or whether there might be a red spotlight shining on the table.

Participants were assigned to either one of the stories and asked to what extent “they agreed or disagreed with the claim that John knows that the table is red” (101–102). It turns out that they were more willing to attribute knowledge in *Furniture-Plain* than in *Furniture-Error*. This is puzzling on the doxastic account. After all, the possibility of abnormal lighting conditions is made salient only to the readers of the story, not the protagonist within the story. So why should the protagonist's doxastic state be affected?

Doxasticists can respond to this challenge by appeal to egocentric bias, a general tendency to project our own mental states onto others.<sup>19</sup> To accommodate the data from ignorant stakes cases, they can appeal, more specifically, to *epistemic* egocentrism, a pervasive failure to suppress one's own privileged knowledge when assessing others. Royzman et al. (2003: 38), for instance, 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 can use this bias to explain intuitions about ignorant stakes cases as follows. Participants in e.g. the *Ignorant-Typo-High* condition know that the stakes are high for John because they've read about this. Given epistemic egocentrism, they cannot suppress this privileged knowledge and project it onto him, despite the stipulation to the contrary. Once the protagonist is *assumed* to know about his high stakes, STAKES may lead people to conclude that he lacks the belief required for knowledge, etc.

To accommodate the data from cases where the error-possibility is presented only to the study participants, a slightly different form of egocentric bias needs to be invoked,

<sup>19</sup> See Nagel (2008: 292–293), (2010a: 426–427) for this idea in the context of stakes effects, and Nagel (2010b: 303) and Nagel and Smith (2017: 100) for this idea in the context salient alternative effects.

what we could call *attentional* egocentrism. Gilovich et al. (2000: 212), for instance, point out that “it might be easy to confuse how salient something is to oneself with how salient it is to others.” To illustrate, participants in one of their studies 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. Something similar might happen in studies where an error-possibility is salient only to the study participants. They egocentrically assume that the possibility is also salient to the protagonist of the story. Now SALIENCE can be used as before to explain salient alternative effects.<sup>20</sup>

Doxasticists endorse an error-theory when it comes to ignorant stakes cases and cases where the error-possibility is salient only to the study participant. On their view, participants *mistakenly* assume that the protagonists of the stories share their knowledge and their concerns due to egocentric bias. This contrasts with the doxasticist account of basic cases and cases with a stipulation of belief. In these cases too, participants make certain assumptions about the protagonists according to doxasticism, namely, that they lack relevant beliefs or that they are rushed. These assumptions, however, supposedly follow from the perfectly justified and presumably correct assumption that the cases are intended to be normal. As such, they can hardly be called a mistake.

## 5 Objections

This completes my presentation of what I take to be a promising version of doxasticism. Stated in this way, a number of worries that have been raised in the literature no longer apply.<sup>21</sup> Other worries remain, and I will point them out in what follows.<sup>22</sup> I will

<sup>20</sup> An analogous story can be told to accommodate salient alternative effects in third-person cases as confirmed in Grindrod et al. (2019). Stakes effects in third-person cases would be more troublesome. For even if we project what we know and what is salient to us onto others, it doesn't follow that we project the high stakes of a speaker in a given story onto all the other protagonists in the story. See Dinges (2018a: §4). Stakes effects in third-person cases haven't been confirmed experimentally though, so doxasticists needn't be too worried yet.

<sup>21</sup> Sripada and Stanley (2012: 20–21), for instance, worry that doxasticism collapses into anti-intellectualism. We've seen how these views differ. Pinillos (2012: 202–203) worries about ignorant stakes cases and cases where belief is stipulated. These cases can be addressed as described above. Stoutenburg (2016: 2037–2039) and Gerken (2017: 287) raise a concern that ignores the principle SALIENCE as an integral part of the doxastic account (see Dimmock 2019: 10n23 for a related observation). Shin (2014: 174) raises a concern that ignores the tracking assumption and its motivation (footnote 12). He raises some further challenges (174–177). The discussion in §4.3 provides ammunition to respond. I have also provided a response to Pynn's (2014: 130) concern from “in a position to know.”

<sup>22</sup> Dimmock (2019) raises additional concerns. First, he worries that epistemic egocentrism should lead participants to project their knowledge e.g. of the fact that the bank will be open onto the protagonists in the cases and thus lead them to ascribe knowledge across the board (3425–3427). But as he himself notes (Sect. 4), egocentric bias is only partial. Thus, participants should only become somewhat more inclined to ascribe knowledge. This is compatible with the fact that their inclinations to ascribe knowledge differ between conditions and are lower in e.g. *Furniture-Error* than in *Furniture-Plain* due to the (egocentrically

address the doxastic account of salient alternative effects and the doxastic account of stakes effects separately because they give rise to different concerns.

### 5.1 Doxasticism and salient alternative effects

One set of concerns with the doxastic account of salient alternative effects arises from the appeal to egocentric bias in cases where the error-possibility is presented only to the study participants. This appeal to egocentric bias generates at least two problematic predictions.

First, egocentric bias can be reduced by strengthening people's desire to give correct answers. Epley et al. (2004: 332), for instance, suggest that egocentric bias can be reduced through monetary rewards for accuracy. Readers of *Furniture-Error* should thus be less inclined to project their concern with an error-possibility onto the protagonist of the story when they have adequate incentives. Given the doxastic account, they should thus become more inclined to ascribe knowledge, for they supposedly deny knowledge only because they project their concerns. This prediction has been disconfirmed, however. Alexander et al. (2014: 108–110) presented participants with *Furniture-Error*. They grouped participants in terms of their “need for cognition” (NFC), a concept from cognitive psychology, which tracks people's “intrinsic motivation to give a great deal of care and attention to cognitive tasks” (108). High NFC naturally comes with a stronger motivation to give correct answers. It should thus reduce egocentric projection and, given the egocentric bias account, increase our willingness to ascribe knowledge. In fact, however, NFC had no effect on participants' responses.

Second, egocentric bias is only partial (Dimmock 2019: Sect. 4). People project their concerns onto others, but they don't assume that everybody is psychologically exactly alike. With respect to the T-shirt study, for instance, Gilovich et al. (2000: 214) point out that, “[a]lthough the target participants overestimated the salience of their T-shirt, their estimates were nonetheless grounded in reality.” We would thus expect error-possibility effects to be stronger when the error-possibility in question is stipulated to be presented to the subject, compared to when it is only mentioned in the case description and allegedly leads to error-possibility effects via egocentric bias. Alexander et al. (2014: 103–105) find no such correlation. They compared responses to *Furniture-Error* with responses to otherwise similar cases where the error-possibility is mentioned by a protagonist of the case, or where the protagonist is described as thinking about the error-possibility without mentioning it aloud. Agreement ratings with the relevant knowledge utterance were entirely unaffected by these modifications.<sup>23</sup>

In summary, egocentric bias doesn't seem to be driving our intuitions about cases where the error-possibility in question is salient only to the study participants. Even if the other assumptions in doxasticism are correct, this shows that the view is insufficiently general, leaving important data unexplained.

Footnote 22 continued

projected) salience of an error-possibility. Dimmock raises a second, stronger concern (3427–3429), that may complement my criticism below.

<sup>23</sup> Since Alexander et al. (2014: 105) fail to appreciate the partiality of egocentric bias, they mistakenly take these data to support the egocentric account.



Another concern with doxasticism arises from sentences of the following kind.

**CF1** Hannah in *Bank-Plain* knows that the bank will be open. She wouldn't know this though if Sarah had mentioned the possibility of changed opening hours.

The counterfactual in CF1 sounds odd if not outright false (e.g. Hawthorne 2004: 166; Dimmock 2018: 2005–2006). Given SALIENCE, and the assumption that knowledge entails belief, it should come out as true though. For given these principles, a shift in salient alternatives may shift our beliefs and thereby what we know. Moreover, if, by the previously mentioned tracking assumption, we track this prediction when assessing fictitious cases such as the bank cases, then the counterfactual in CF1 should not only be true; we should also judge it to be true. After all, assessing a counterfactual seems very similar to assessing a fictitious case where the antecedent conditions hold. This makes it deeply puzzling why the counterfactual in CF1 sounds false.

Dimmock (2019) suggests a way for doxasticists to handle counterfactuals such as those in CF1. He notes the oddity of “[John] doesn't know the table is red. But if his son hadn't raised the possibility that the table is white but illuminated by red lights, he would know that it's red” (3414). He claims that doxasticists can explain this oddity as follows (3417). The possibility of abnormal lighting conditions is salient to us as readers. Via egocentric bias, we assume that it is also salient to John when we counterfactually consider him in a situation where the error-possibility isn't salient to him. Thus, John appears to lack knowledge in this counterfactual situation, and the indicated counterfactual comes out as false, in line with our intuitions. Unfortunately, an analogous story cannot be told about the counterfactual in CF1, which is an inverse of Dimmock's counterfactual. Here we counterfactually consider a situation where the error-possibility *is* salient to Hannah, and so Hannah *lacks* knowledge according to doxasticism. Hannah will continue to lack knowledge in this situation, and thus the counterfactual will remain true, even if we egocentrically project our concerns onto her.

One might suggest that CF1 as a whole sounds bad due to the initial, categorical sentence. In particular, one might suggest that we assess this sentence retrospectively, after reading the conditional and hence after the possibility of changed opening hours has become salient to us. Now the initial sentence may appear false because Hannah may appear to lack knowledge once we egocentrically assume that she's concerned with this error-possibility too. First, though, it's the counterfactual in CF1 that sounds false, not the initial sentence. Second, the following text still sounds bad even though it doesn't say (or suggest) that Hannah knows the bank will be open: “Hannah in *Bank-Plain* may or may not know that the bank will be open. But she wouldn't know this if Sarah had mentioned the possibility of changed opening hours.”

The principles of doxasticism jointly give rise to faulty predictions about CF1. Which of these principles should we reject? For what it's worth, I think SALIENCE is well motivated, and I think knowledge entails belief. I therefore think that CF1 can be true. It still doesn't follow that CF1 should *sound* true, which is the problematic prediction. We need the tracking assumption, and I think this assumption is false, as the intuitive oddity of CF1 brings out.



To independently motivate the rejection of the tracking assumption, notice that readers of e.g. the bank cases will focus on the evidential requirements on knowledge because the evidence available to the protagonist takes center stage in the case descriptions. Meanwhile, the doxastic dimension of knowledge is completely out of focus. It is therefore natural to expect that shifts in this latter dimension will be treated as irrelevant if registered at all and that participants will reach a judgement about knowledge (or “knowledge”) as soon as they have settled whether the evidence described suffices for knowledge (or “knowledge”).<sup>24</sup> This would also explain why intuitions about CF1 aren’t entirely stable and why they can reverse as soon as we consider views like doxasticism. Once we consider such views, we focus on the doxastic dimension of knowledge and therefore realize that CF1 isn’t all that problematic after all.

It may be instructive here to contrast doxasticism with my own favored account of the data (see Dinges 2018b for details; and Buckwalter 2019 for confirmation of a key prediction). This account starts out from the assumption that salient error-possibilities lead us to overestimate the probability of error, and it uses this assumption to explain salient alternative effects as follows. Consider bank case studies from the perspective of a study participant. Study participants have to assess whether the protagonist knows that the bank will be open and, on natural assumptions about the nature of knowledge, they will thereby have to assess the evidential probability of this proposition on the protagonist’s evidence. Specifically, they have to assess how probable it is that the bank will be open given the protagonist’s memory of a visit at the bank two weeks before. By the above assumption, mentioning an error-possibility leads to lower estimates here and, consequently, to denials of knowledge.

On this account, it is unsurprising that stipulating belief is ineffective, that we find salient alternative effects for “in a position to know” and that it doesn’t matter whether the error-possibility is described as being salient to the protagonist of the story. Study participants deny knowledge not because of a lack of belief on the part of the protagonist but because they themselves deem the evidence too weak. And they deem the evidence too weak because an error-possibility has been brought to *their* attention. Additionally, the account no longer predicts that CF1 sounds true, as long as we focus on the evidential rather than the doxastic dimension of knowledge. On the suggested account, salient error-possibilities shift our *estimates* of evidential probability. They don’t shift the evidential probability itself. The counterfactual in CF1 is therefore false as far the evidential dimension of knowledge is concerned, and it sounds false for this reason.

This account is an error-theory, unlike doxasticism (egocentric bias aside). Denials of knowledge in the face of salient alternatives result from an overestimation of the probability of error. Relatedly, the account focuses on the psychology of the study participant rather than the psychology of the protagonists of the story, unlike doxasticism (egocentric bias aside). In particular, it focuses on the estimates participants

---

<sup>24</sup> See Gerken 2017: 116–119 for empirically informed discussion of the claim that we often form epistemic judgements based on a limited set of factors that happen to be in focus. The suggested account is consistent with the previously mentioned data that people are sensitive to the influence of non-traditional factors when they assess other people’s beliefs (footnote 12). For when assessing beliefs, you automatically focus on doxastic factors.

make about how probable the target proposition is, entailing that these estimates are systematically biased when error-possibilities are raised to salience.

The suggested view still bears an important relation to doxasticism. While SALIENCE doesn't feature in the account, it still comes out as an immediate prediction. If salient alternatives lead us to overestimate the probability of error, as the suggested view would have it, they should also make us less certain or altogether lose belief. After all, certainty and belief presumably track estimated probability. SALIENCE bears this out.

## 5.2 Doxasticism and stakes effects

Let's turn to the doxastic account of stakes effects. Notice to begin with that the appeal to egocentric bias comes out as less problematic here than in the case of salient alternative effects. This is primarily because immediately relevant data are unavailable. Additionally, the available data tentatively confirm the respective predictions. If the egocentrism account was correct, one would expect the readers of *Typo-High* to be more hesitant to ascribe knowledge than the readers of *Ignorant-Typo-High*. After all, and as indicated, egocentrism is only partial in that we retain some sense of the difference between our mental states and other people's mental states. This is borne out by the data. As Pinillos (2012: 21–22) points out, “answers to *Ignorant-Typo-High* (Median = 3) and *Typo-High* (Median = 5) were different,” and the difference goes in the right direction.

As before, though, sentences of the following kind put pressure on the doxastic account.

CF2 Peter in *Typo-Low* needs to read his paper two times before he knows there are no typos. If more was at stake, though, more reads would be required for knowledge.

The counterfactual in CF2 sounds false just like the counterfactual in CF1 does.<sup>25</sup> But given the tracking assumption, it's difficult to see how doxasticists can make sense of this. We supposedly realize that protagonists in high stakes cases normally need more evidence before they form a belief and thus obtain knowledge. Why should we suddenly fail to realize this when assessing counterfactuals such as the one in CF2?

Nagel (2010a: 426) suggests an insufficiently general way to handle such counterfactuals. She considers the sentence “He knows, but he wouldn't have known if more had been at stake,” which sounds odd too. To explain this, she notes “that we ordinarily expect an ordinary rational subject who is shifted into higher stakes to feel increased epistemic anxiety and do the evidence-collecting work appropriate to his increased anxiety.” Thus, if the stakes had been higher, the relevant subject would still have had knowledge, and Nagel's counterfactual comes out as false, as it should. Unfortunately, a similar story cannot be told about the counterfactual in CF2, which isn't directly about knowledge but about the (normal) evidential requirements for

<sup>25</sup> For related observations, see e.g. DeRose (2009: 194), Fantl and McGrath (2009: 208–212), Hawthorne (2004: 177n40), Schaffer (2006: 104n6), Stanley (2005: 106), Blome-Tillmann (2009), and Dimmock (2018).

knowledge. Even if we ordinarily do enough to meet these requirement, they still shift according to the doxastic account and thus the counterfactual in CF2 still comes out as true, contrary to our intuitions.

What is wrong with doxasticism? My diagnosis here is the same as before in the case of salient alternative effects. The faulty assumption is the tracking assumption. Knowledge entails belief and a principle like STAKES governs belief, so CF2 can be true. The oddity of CF2, however, brings out that people don't track this prediction, at least as long as they consider familiar cases like the typo cases. This is because the evidential situation of the protagonist takes center stage in these cases, which leads participants to focus on the evidential dimension of knowledge to the exclusion of the doxastic dimension. As before, this helps to explain why CF2 can sound true once we shift our focus to views like doxasticism. This shifts our focus to the doxastic dimension of knowledge, which is sensitive to stakes.

Once more, it might be instructive to contrast the doxastic account of stakes effects with my own favored account of the data (see Dinges, ms for details). This account starts from the idea that high stakes tend to make it very costly to overestimate the evidential probability of the target proposition. In *Typo-High*, for instance, it would be very costly to overestimate the evidential probability of the proposition that no typos remain after any given number of rounds of proofreading; for this could lead to premature submission with all the devastating consequences described in the story. Participants run with lower estimates in order to make sure that if they err, they err on the safe side. Consequently, they are less inclined to ascribe knowledge.

On this view, it is unproblematic to explain why stipulating belief or shifting to “in a position to know” doesn't affect stakes effects and why these effects remain when the protagonist is ignorant of what is at stake. Participants deny knowledge because they deem the protagonist's evidence too weak, not because they think he lacks a relevant belief. And they deem the evidence too weak as long as *they* know about the potential costs of overestimation. CF2 isn't predicted to sound true either if we focus exclusively on the evidential dimension of knowledge. On the suggested account, stakes affect *estimates* of evidential probability, but they don't affect evidential probability itself. Thus, the counterfactual in CF2 is false as far as the evidential dimension of knowledge is concerned, which is why it sounds false.

As before, the suggested theory is an error-theory that focuses on the psychology of the study participants rather than the protagonists of the cases, unlike doxasticism (egocentric bias aside). On the suggested view, stakes lead study participants to underestimate the evidential probability of the target proposition, placing overly demanding constraints on knowledge.

A connection to doxasticism remains. STAKES plays no role in this theory, but, just like SALIENCE before, this principle comes out as an immediate prediction. If high stakes lead to lower estimates of probability, they will also lead us to report lower levels of certainty, which, after all, track our assessments of probability. The empirical evidence supporting STAKES bears this out.

## 6 Conclusion

Doxasticism is more promising than it is often made out to be. Indeed, most of the worries that have been raised in the literature can be deflected. Still, many direct predictions of this account have been disconfirmed and to the extent that one can make sense of this at all, special pleading is required in each case. The central principles STAKES and SALIENCE may still point into the right direction, suggesting that estimates of evidential probability are affected by non-traditional factors. As indicated, this latter assumption promises to offer a much neater, error-theoretic account of the data, but the details must be treated elsewhere.

**Acknowledgements** I am grateful to Jie Gao, Mikkel Gerken, Julia Zakkou, two anonymous referees, the members of the *Oberseminar* in Erlangen as well as audiences in Amsterdam, Cologne, Hamburg and Osnabrück for very helpful feedback on earlier versions of this paper. Open Access funding provided by Projekt DEAL.

**Funding** The study was funded by Deutsche Forschungsgemeinschaft (Grant Number DI 2172/1-1).

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Alexander, J., Gonnerman, C., & Waterman, J. (2014). Salience and epistemic egocentrism: An empirical study. In J. R. Beebe (Ed.), *Advances in experimental epistemology* (pp. 97–118)., *Advances in experimental philosophy* New York, NY: Bloomsbury Publishing.
- Bach, K. (2005). The emperor's new 'knows'. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy. Knowledge, meaning, and truth* (pp. 51–90). Oxford: Oxford University Press.
- Bach, K. (2008). Applying pragmatics to epistemology. *Philosophical Issues*, 18(1), 68–88.
- Bach, K. (2010). Knowledge in and out of context. In J. K. Campbell, M. Orourke, & H. Silverstein (Eds.), *Knowledge and skepticism* (pp. 105–136)., *Topics in contemporary philosophy* Cambridge, MA: MIT Press.
- Blome-Tillmann, M. (2009). Contextualism, subject-sensitive invariantism, and the interaction of 'knowledge'-ascriptions with modal and temporal operators. *Philosophy and Phenomenological Research*, 79(2), 315–331.
- Blome-Tillmann, M. (2014). *Knowledge and presuppositions*. Oxford: OUP.
- Brown, J. (2008). Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs*, 42(2), 167–189.
- Buckwalter, W. (2014). The mystery of stakes and error in ascriber intuitions. In J. R. Beebe (Ed.), *Advances in experimental epistemology* (pp. 145–173)., *Advances in experimental philosophy* New York, NY: Bloomsbury Publishing.
- Buckwalter, W. (2019). Error possibility, contextualism, and bias. *Synthese*. <https://doi.org/10.1007/s11229-019-02221-w>.
- Buckwalter, W., & Schaffer, J. (2015). Knowledge, stakes, and mistakes. *Noûs*, 49(2), 201–234.
- Dafoe, A., Zhang, B., & Caughey, D. (2018). Information equivalence in survey experiments. *Political Analysis*, 26(4), 399–416.

- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52(4), 913–929.
- DeRose, K. (2009). *The case for contextualism.*, Knowledge, skepticism, and context, 1 Oxford: OUP.
- Dimmock, P. (2018). Strange-but-true. A (quick) new argument for contextualism about ‘know’. *Philosophical Studies*, 175(8), 2005–2015.
- Dimmock, P. (2019). Knowledge, belief, and egocentric bias. *Synthese*, 196(8), 3409–3432.
- Dinges, A. (ms). Better safe than sorry. Explaining stakes effects on knowledge attributions.
- Dinges, A. (2018a). Anti-intellectualism, egocentrism and bank case intuitions. *Philosophical Studies*, 175(11), 2841–2857.
- Dinges, A. (2018b). Knowledge and availability. *Philosophical Psychology*, 31(4), 554–573.
- Dinges, A. & Zakkou J. (forthcoming). Much at stake in knowledge. *Mind & Language*.
- Epley, N., Keysar, B., van Boven, L., et al. (2004). Perspective taking as egocentric anchoring and adjustment. *Journal of Personality and Social Psychology*, 87(3), 327–339.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford: OUP.
- Farkas, K. (2015). Belief may not be a necessary condition for knowledge. *Erkenntnis*, 80(1), 185–200.
- Francis, K., Beaman, P., & Hansen, N. (2019). Stakes, scales, and skepticism. *Ergo*, 6(16), 427–487.
- Ganson, D. (2008). Evidentialism and pragmatic constraints on outright belief. *Philosophical Studies*, 139(3), 441–458.
- Gao, J. (2019). Credal pragmatism. *Philosophical Studies*, 176(6), 1595–1617.
- Gerken, M. (2017). *On folk epistemology: How we think and talk about knowledge*. Oxford: OUP.
- Gerken, M., Gonnerman, C., Alexander, J., et al. (2020). Salient alternatives in perspective. *Australasian Journal of Philosophy*. <https://doi.org/10.1080/00048402.2019.1698625>.
- 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.
- Grindrod, J., Andow, J., & Hansen, N. (2019). Third-person knowledge ascriptions: A crucial experiment for contextualism. *Mind and Language*, 34(2), 158–182.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: OUP.
- Hawthorne, J., Rothschild, D., & Spectre, L. (2016). Belief is weak. *Philosophical Studies*, 173(5), 1393–1404.
- Ichikawa, J. J., & Steup, M. (2014). The analysis of knowledge. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy*. Stanford, CA: The Metaphysics Research Lab: Center for the Study of Language and Information: Stanford University.
- Kelley, H. H. (1972). Attribution in social interaction. In E. Jones, D. Kanouse, H. H. Kelley, et al. (Eds.), *Attribution perceiving the causes of behavior* (pp. 1–26). Morristown, NJ: General Learning Press.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108(3), 480–498.
- Mayselless, O., & Kruglanski, A. W. (1987). What makes you so sure? Effects of epistemic motivations on judgmental confidence. *Organizational Behavior and Human Decision Processes*, 39(2), 162–183.
- Myers-Schulz, B., & Schwitzgebel, E. (2013). Knowing that P without believing that P. *Noûs*, 47(2), 371–384.
- 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. *The Philosophical Quarterly*, 60(239), 286–306.
- Nagel, J. (2012). The attitude of knowledge. *Philosophy and Phenomenological Research*, 84(3), 678–685.
- Nagel, J., Juan, V. S., & Mar, R. A. (2013). Lay denial of knowledge for justified true beliefs. *Cognition*, 129(3), 652–661.
- Nagel, J., & Smith, J. J. (2017). The psychological context of contextualism. In J. J. Ichikawa (Ed.), *The routledge handbook of epistemic contextualism* (pp. 94–104)., Routledge handbooks in philosophy New York, NY: Routledge.
- 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. R. Beebe (Ed.), *Advances in experimental epistemology* (pp. 9–44)., Advances in experimental philosophy New York, NY: Bloomsbury Publishing.
- Pynn, G. (2014). Unassertability and the appearance of ignorance. *Episteme*, 11(02), 125–143.

- Radford, C. (1966). Knowledge. By examples. *Analysis*, 27(1), 1–11.
- Rose, D., Machery, E., Stich, S., et al. (2019). Nothing at stake in knowledge. *Noûs*, 53(1), 224–247.
- Ross, J., & Schroeder, M. (2014). Belief, credence, and pragmatic encroachment. *Philosophy and Phenomenological Research*, 88(2), 259–288.
- 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.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Noûs*, 35(4), 477–514.
- 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(02), 157–180.
- Sripada, C. S., & Stanley, J. (2012). Empirical tests of interest-relative invariantism. *Episteme*, 9(1), 3–26.
- Stanley, J. (2005). *Knowledge and practical interests.*, Short philosophical books Oxford: OUP.
- Stoutenburg, G. (2016). Strict moderate invariantism and knowledge-denials. *Philosophical Studies*, 174(8), 2029–2044.
- Tebben, N. (2019). Knowledge requires commitment (instead of belief). *Philosophical Studies*, 176(2), 321–338.
- Vogel, J. (1990). Are there counterexamples to the closure principle? In M. D. Roth & G. Ross (Eds.), *Doubting: Contemporary perspectives on skepticism* (pp. 13–27). Dordrecht: Kluwer Academic Publishers.
- Weatherson, B. (2005). Can we do without pragmatic encroachment? *Philosophical Perspectives*, 19(1), 417–443.
- Weatherson, B. (2017). Interest-relative invariantism. In J. J. Ichikawa (Ed.), *The routledge handbook of epistemic contextualism* (pp. 240–254)., Routledge handbooks in philosophy New York, NY: Routledge.
- Williamson, T. (2005). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *The Philosophical Quarterly*, 55(219), 213–235.
- Wright, C. (2017). The variability of 'knows'. An opinionated overview. In J. J. Ichikawa (Ed.), *The routledge handbook of epistemic contextualism* (pp. 13–31)., Routledge handbooks in philosophy New York, NY: Routledge.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.