

Knowledge, Belief, and Egocentric Bias

Paul Dimmock

Abstract

Changes in conversationally salient error possibilities, and/or changes in stakes, appear to generate shifts in our judgments regarding the correct application of ‘know’. One prominent response to these shifts is to argue that they arise due to shifts in belief and do not pose a problem for traditional semantic or metaphysical accounts of knowledge (or ‘know’). Such doxastic accounts face familiar difficulties with cases where knowledge is ascribed to subjects in different practical or conversational situations from the speaker. Jennifer Nagel has recently offered an ingenious response to these problematic cases—appeal to egocentric bias. Appeal to this kind of bias also has the potential for interesting application in other philosophical arenas, including discussions of epistemic modals. In this paper, I draw on relevant empirical literature to clarify the nature of egocentric bias as it manifests in children and adults, and argue that appeal to egocentric bias is ill-suited to respond to the problem cases for doxastic accounts. Our discussion also has significant impact on the prospects for application of egocentric bias in other arenas.

Keywords: Psychological Bias; Epistemic Egocentrism; Knowledge Ascriptions; Classical Invariantism

Introduction

As is familiar from examples involving banks, airports, and painted mules, changes in conversationally salient error-possibilities, and/or changes in how much is at stake, appear to generate shifts in our judgments regarding the correct application of ‘know’.¹ *Contextualists* have argued that these judgments motivate the semantic thesis that the contents of ‘knowledge’-ascribing and denying sentences vary with context. *Impurists* have argued instead that these judgments motivate the metaphysical thesis that knowledge is constitutively tied to practical or conversational factors.² But more conservative theorists have rejected the claim that the shifts in our judgments about ‘know’ warrant our embrace of surprising metaphysical or semantic theses about knowledge (or ‘know’). Among these more conservative theorists are those who suggest that our judgments are an upshot of the fact that shifts in stakes and/or salient error possibilities lead to shifts in the presence of belief (see esp. Bach 2005: §V & Nagel 2008, 2010a, 2010b, 2011).³

Such belief-centric (or ‘doxastic’) proposals face a number of familiar challenges and difficulties. Some prominent obstacles include accounting for examples where belief is stipulated to be present and for examples where knowledge is ascribed (or denied) to a subject in a different practical or conversational situation from the speaker. In a series of papers, Jennifer Nagel (2008, 2010a, 2010b, 2011) has attempted to address these challenges for doxastic proposals. Although Nagel’s proposals have not garnered much

¹ See e.g. Vogel 1990: 15-16, Cohen 1999: 58, DeRose 2009: 1-5, Lewis 1996. There is now a substantial empirical literature investigating the extent to which these shifts in judgment are exhibited by ordinary speakers. There seems to be some good evidence that shifts in the salience of error possibilities generate shifts in ordinary speakers’ judgments regarding ‘know’, but the situation is arguably less clear in regard to stakes effects. See e.g. Buckwalter 2014, Schaffer & Knobe 2012, and Buckwalter & Schaffer 2014 for some relevant empirical work and discussion.

² Prominent contextualist accounts include Cohen 1999, DeRose 1995, 2009, Lewis 1996, Blome-Tillmann 2014, Ichikawa 2017. Impurist (or ‘anti-intellectualist’) accounts include Hawthorne 2004: ch. 4, Stanley 2005, Fantl & McGrath 2009, Weatherson 2005, 2017.

³ It is possible to pursue a doxastic approach to explaining shifts in our judgments about ‘know’ that is metaphysically or semantically non-conservative (see e.g. Weatherson 2005). In the present paper, I shall focus on conservative (i.e. classical invariantist) attempts to pursue a doxastic approach, but our discussion plausibly has significance for some non-conservative doxastic approaches as well.

serious attention in the literature, her attempt to revive the doxastic approach seems especially worthy of consideration in light of recent criticism of other attempts by conservatively-minded theorists to explain the aforementioned shifts in our judgments about ‘know’ (see e.g. Nagel 2010a: 286-301, Blome-Tillmann 2013, Dimmock & Huvenes 2014, Dinges forthcoming (a)). It also seems worthy of consideration in light of an interesting and innovative aspect of Nagel’s proposals—her appeal to egocentric bias. As we shall see later on, appeal to egocentric bias would seem to have the potential to address problems both for various epistemological theories, and also for theories in a range of other philosophical arenas, including accounts of epistemic modals, predicates of personal taste, and moral claims. A more careful look at how egocentric bias functions would therefore seem to be of some significant and general philosophical interest.

§1. A Doxastic Approach

Consider the following examples (adapted from Nagel 2010a: 287; Cohen 2002: 312-3):

Table A. John is in a store looking at what appears to be a bright red table a few yards ahead of him. John’s young son asks John, ‘Do you know that the table is red?’. John replies, ‘Yes, I know it’s red’.

Table B. John is in a store looking at what appears to be a bright red table a few yards ahead of him. John’s young son asks John, ‘Do you know that the table is red?’, and remarks that the table would appear just the same to John if it was white but illuminated by red lights. John replies, ‘No, I don’t know that it’s red’.

Assume that in both *Table A* and *Table B* the table is indeed red, and the lighting conditions normal. The standard story, applied to our examples, is that John’s knowledge ascription

seems true in *Table A* and his knowledge denial seems true in *Table B*.⁴ *Non-sceptical* epistemologists concede that John's knowledge ascription in examples like *Table A* is true. These theorists therefore have no trouble accommodating the judgment that his knowledge ascription in *Table A* seems true. But John's grounds for believing that the table is red seem to be the same in *Table A* and *Table B*. It might therefore seem that if John knows in *Table A*, he knows in *Table B*. The challenge for non-sceptical epistemologists is therefore to explain why it nevertheless seems true for John to *deny* that he knows in examples like *Table B*.

As noted in the Introduction, some prominent non-sceptical epistemologists have argued that our judgments about examples like *Table B* can be explained if we accept either contextualism or impurism. But more conservative non-sceptical epistemologists—*classical invariantists*—reject these contentious theories, and maintain that 'know' is not semantically context sensitive and that knowledge is not constitutively tied to 'non-epistemic' factors, such as salient error possibilities or stakes.⁵ In order to defend their position, classical invariantists therefore require some alternative explanation for our judgments regarding examples like *Table B*.

One natural classical invariantist suggestion is that the reason John's knowledge denial seems true in *Table B* is that considering the possibility that the table is white but illuminated by red lights causes John to lose his belief that the table is red. If that's right, then—assuming belief is required for knowledge—it follows that John's utterance of 'I

⁴ For some relevant empirical work on these kinds of judgments, see e.g. Schaffer & Knobe 2012; Buckwalter 2014. Note that to generate the reported judgments it may be necessary to amend *Table B* to ensure that the possibility of tricky lighting becomes sufficiently salient (see Schaffer & Knobe 2012: 19-22). If necessary, the discussion to follow could be recast in terms of such amended cases.

⁵ For characterisation of the relevant positions in the debate, see e.g. DeRose (2009: 1-49) or MacFarlane (2014: ch. 7). Note that classical invariantists also reject the semantic claim associated with relativist or perspectival accounts of 'know'—viz. the claim that the contents of 'knowledge'-ascribing and denying sentences are only true relative to some additional 'epistemic standards' parameter (see e.g. MacFarlane 2005, 2014: ch. 7).

don't know that the table is red' in *Table B* is in fact true. It is therefore unsurprising that it seems true.

There seem to be three broad mechanisms via which John might lose his belief in *Table B*:⁶

Lowers credence. Some empirical research suggests that considering additional error possibilities leads to lower levels of subjective confidence (Kelley 1972). It might therefore be proposed that considering the possibility of tricky lighting causes John to lower his credence (degree of confidence) in the proposition that the table is red such that his credence is no longer high enough for him to count as believing that it is red (cf. Nagel 2010b: 422).

Raises threshold. Some theorists have suggested that the level of credence required for belief is determined (in part) by the practical or conversational situation of the subject (Weatherson 2005). It might therefore be proposed that John's consideration of the possibility of tricky lighting raises the credence threshold John must meet to count as believing that the table is red such that, even if John does not lower his credence upon considering the possibility of tricky lighting, he nevertheless ceases to believe that the table is red.⁷

Other factors. Having a belief might not simply be a matter of having a credence above a certain threshold—it might require something else in addition (or instead). It might

⁶ These three broad approaches can be found in Nagel (2010b), though Nagel does not take pains to distinguish them. Bach (2005: §V) does not indicate the mechanism via which a subject like John might lose his belief, merely remarking (in regard to a similar case) that the subject's belief "may be shaken somewhat".

⁷ Note that this particular account of how John loses his belief may ultimately require accepting impurism, and thus be unacceptable to those seeking to defend classical invariantism. See Weatherson (2005) and Nagel (2010b: 417-8) for some relevant discussion.

therefore be proposed that considering the possibility of tricky lighting removes John's belief that the table is red by impacting that 'something else'.

For example, it might be proposed that believing that P requires—perhaps in addition to having a certain credence in P—the kind of psychological conviction associated with 'taking it to be settled' that P or 'having one's mind made up' that P (cf. Nagel 2010b: 416-21). It might then be suggested that in situations where we are considering ways we might be mistaken with respect to P, we often require additional evidence before being psychologically able to (e.g.) take it to be settled that P.⁸ In that case, considering the possibility that the table is white but illuminated by red lights might remove John's belief that the table is red *not* because it lowers his credence that the table is red, and not because the credence threshold for belief goes up, but rather because considering that possibility causes John to no longer take it to be settled that the table is red. Call this view *removes psychological conviction*.⁹

The particular mechanism via which John loses his belief is not central to the discussion to come, so I shall remain neutral on that issue in what follows.¹⁰ In addition, I shall continue to focus on examples, like *Table A* and *Table B*, that concern how shifts in salient error possibilities impact our judgments about the correct application of 'know'. Doxastic approaches to examples involving shifts in practical factors, such as stakes, meet with parallel problems, and the discussion to follow can be fairly straightforwardly extended to

⁸ See Nagel (2010b: esp. 416-21, 2011: 13-15, 2010a: 303) for development of ideas along similar lines, and discussion of relevant psychological literature.

⁹ It may be natural to pursue a similar proposal if one thinks that the doxastic requirement on knowledge is not belief, but is rather 'being sure' or 'being (subjectively) certain' (see e.g. DeRose 2009: 186n).

¹⁰ Nagel appears to show preference for a view along the lines of *removes psychological conviction* (see e.g. Nagel 2010b: 418). One possible advantage of this proposal is that it may be more naturally suited to explaining our judgments regarding just *how much* additional evidence John requires in order to know that the table is red (see esp. Nagel 2011: 13-5; also 2010a: 303). Issues surrounding how much additional evidence subjects like John need to possess in order to know will return later (§2, §6).

cover such examples. To keep things manageable, I shall therefore largely ignore stakes-based cases in what follows.

§2. Challenges

An immediate concern for those sympathetic to the kind of doxastic approach sketched in the previous section is that we exhibit similar judgments about the correct application of ‘know’ even when belief is stipulated to be present. For example, we can imagine a case, *Table B**, which is just like *Table B* except that *Table B** contains the additional stipulation that John believes, on the basis of how the table looks, that the table is red. The concern is that even with such an additional stipulation in place, it seems true for John to respond, ‘I don’t know that the table is red’ (cf. Nagel 2010a: 287-8, DeRose 2009: 1-2). How is this to be squared with a doxastic approach to explaining the shifts in our judgments about ‘know’?

Nagel (2008, 2010a, 2010b, 2011) puts forward a novel strategy for responding to such ‘stipulated-belief’ cases. In regard to a case similar to *Table B**, she writes:

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. (Nagel 2010a: 303)

Roughly, then, the strategy is that if it is stipulated that John believes that the table is red despite actively considering the possibility that it is white but illuminated by red lights (and

despite not checking for red lights), he will appear to have formed his belief via the influence of some epistemically problematic factor, such as wishful thinking. The presence of such a factor would seem to us to render John's belief formation insufficiently accurate (reliable), and so we judge that John lacks knowledge.¹¹

A great deal more could be said about this strategy for handling stipulated-belief cases, but for the purposes of the present paper, I shall just grant that it is successful.¹² In what follows, I wish to focus on a further set of problematic cases for the doxastic approach, concerning situations where the subject of a knowledge claim is in a different conversational or practical situation from the speaker. Various cases illustrate the problem. The following are two representative examples:

Table (3rd Person). Suppose that John and his son are in a situation like *Table B*—John's son has just raised the possibility that the table is white but illuminated by red lights, and John utters 'I don't know that the table is red'. But now suppose that John's wife, Alice, is a few feet away from John, with a similar clear view of the table as him. Alice has clearly not overheard John and her son's conversation, and is not considering ways she might be mistaken, such as that the table is white but illuminated by red lights. Alice believes, on the basis of how the table looks, that the table is red.

Suppose that after John issues his knowledge denial, John's son notices his mother looking at the table. John's son says to John that the table would also appear just the same to his mother if it was white but illuminated by red lights, and asks John if his mother knows that the table is red. John responds 'No, she has the same evidence as me. She doesn't know that it's red either.'

¹¹ Nagel (2010b: 420n) offers a slightly different suggestion: that the subject's belief may appear to fall short of knowledge because it will seem to lack the epistemic virtues necessary for knowledge. The differences between this proposal and the one in the main text are not important for the discussion to follow.

¹² Sripada & Stanley (2012: 18-23) criticise Nagel's strategy for handling stipulated-belief cases as either implausible or committed to impurism (and so unsuitable for preserving classical invariantism). Nagel addresses some concerns along these lines in (Nagel 2010b: 427-8; see also 2011: 13-5). Shin (2014: 173-7) also raises some concerns for Nagel's proposal.

In regard to similar cases, it is standardly reported that John's response seems true (see e.g. Nagel 2010a: 287-8, Vogel 1990: 15-16, Cohen 1999, DeRose 2009: 3-6).¹³

Table (Modal Contrast). Suppose that John and his son are in a situation like *Table B*—John's son has just raised the possibility that the table is white but illuminated by red lights, and John utters 'I don't know that the table is red'.

Two shop assistants, Rick and Mona, are standing next to the table, and have overheard John and his son's conversation. Rick and Mona look up and check that there is no red lighting. Rick then asks Mona if she agrees with John that he doesn't know that the table is red, reiterating John's son's observation that the table would appear just the same to John if it was white but illuminated by red lights. Mona replies, 'I agree—he 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'.

In regard to similar cases, it is standardly reported that Mona's response seems strange (see e.g. Hawthorne 2004: 177n, Nagel 2010b: 426, Blome-Tillmann 2009: 320).¹⁴

A doxastic approach, even one supplemented by Nagel's strategy for handling stipulated-belief cases, seems ill-suited to account for our judgments in these cases. Consider *Table (3rd Person)*. In that example, John's son remarks that the table would appear just the same to Alice if it was white but illuminated by red lights, and John utters '[Alice] doesn't know that it's red'. John's utterance seems true, but since Alice is not considering error possibilities, the doxastic approach does not seem to supply any obstacle to Alice knowing

¹³ As in the case of *Table B*, it may be necessary to amend the case to ensure that the possibility of tricky lighting becomes sufficiently salient in order to generate the judgment that John's denial of knowledge to Alice seems true (see fn. 4). (Similar remarks may also apply to *Table (Modal Contrast)*.) Such amendments are not important to the discussion that follows.

¹⁴ Note that Mona's utterance may seem true (and not strange) if we suppose that had John's son not raised the possibility of tricky lighting, John would have looked up at the lighting, and so been able to confirm that the lighting conditions are normal. But I take it that this is not the natural reading of the case. The natural reading is that if his son had not raised the possibility of tricky lighting, John would have had just the same grounds to believe that the table is red that he has in the actual case (roughly, how the table looks), and so would still lack confirmation that that the lighting conditions are normal.

that the table is red. Indeed, as Alice resembles the John-character in *Table A*, it seems that advocates of the doxastic approach (as non-sceptics) should accept that Alice knows that the table is red. So why does John's utterance seem true?

Or consider *Table (Modal Contrast)*. In that example, Rick remarks that the table would seem just the same to John if it was white but illuminated by red lights, and Mona utters '[John] doesn't know that the table is red. But if his son hadn't mentioned the possibility that the table is white but illuminated by red lights, he would know that it's red'. Mona's utterance seems strange, but on the doxastic approach, her utterance is plausibly true. As things actually stand, the possibility of tricky lighting is salient to John. Given the doxastic approach, it should therefore be natural to suppose that John does not believe, and so does not know, that the table is red. However, if John's son had not mentioned the possibility of tricky lighting, the obstacle to John believing that the table is red would presumably be absent, and John would know that the table is red. So why does Mona's utterance seem strange?¹⁵

The general problem underlying cases like *Table (3rd Person)* and *Table (Modal Contrast)* can be usefully stated in terms of the familiar language of 'ruling out' error possibilities. Once an error possibility becomes suitably salient, we are prone to judge as though subjects must be able to rule out that possibility in order to be truly said to 'know' (cf. Lewis 1996). For example, once the possibility that the table is white but illuminated by red lights becomes suitably salient, we are prone to judge as though John and Alice must be able to rule out that possibility—i.e. possess something like the evidence acquired by explicitly looking up and checking the lighting—in order to be truly said to 'know' that the table is

¹⁵ It should be noted that the range of problem cases extends beyond examples like *Table (3rd Person)* and *Table (Modal Contrast)*. Other relevant examples include *Temporal Contrast* cases (see e.g. Stanley 2005: 106), *Third-Person Contrast* cases (see e.g. Neta 2007: 182-3), and *Retraction* cases (see e.g. MacFarlane 2005: §2.3). Our discussion could just as easily have focused on these examples.

red.¹⁶ A positive aspect of the doxastic approach is that it seems suited to explaining why we judge in this way in regard to subjects, like John in *Table B*, who are considering the error possibilities at issue: these subjects need to gather additional evidence in order to naturally form the relevant belief (i.e. to form the relevant belief without the influence of epistemically problematic factors, like wishful thinking). However, we also judge as though those subjects who are *not* considering the error possibilities at issue, such as Alice in *Table (3rd Person)*, must be able to rule out those possibilities in order to be truly said to ‘know’ (cf. Bach 2005: §V). But at least insofar as the error possibilities at issue are distant or improbable ones—such as the possibility that the table is white but illuminated by red lights—subjects who are not considering those error possibilities presumably do not need to be able to rule them out in order to naturally form beliefs.¹⁷ Thus, in regard to examples involving these kinds of subjects, a doxastic approach seems ill-suited to explaining our judgments.¹⁸

§3 Egocentric Bias

The preceding cases, involving attributions/denials of knowledge to subjects in different conversational situations from the speaker, can seem extremely problematic for advocates of the doxastic approach. But Nagel (2010a: 301-6, 2010b: 425-6) provides an ingenious

¹⁶ I take it that we have some intuitive grasp on what is needed for a subject like John to ‘rule out’ that the table is white but illuminated by red lights. Ruling out that possibility seems to require something like the evidence acquired by explicitly looking up and checking the lighting, and something over and above mere statistical evidence for thinking that the relevant tricky-lighting scenario is unlikely. It should be noted that some philosophers think that our intuitive notion of ‘ruling out’ may simply collapse into knowledge that the relevant possibility does not obtain (see e.g. DeRose 1995: 16-7). But this is not important for us here: the present appeal to ruling out error possibilities is being made for illustrative (and not reductive) purposes.

¹⁷ See Nagel (2010b) for extensive discussion of natural (“evidence-based”) belief formation vs. epistemically problematic belief formation.

¹⁸ It might be suggested that the best response to examples like *Table (3rd Person)* and *Table (Modal Contrast)* is to combine a doxastic approach with some other explanatory approach (e.g. the pragmatic approach found in Brown 2006, Rysiew 2007). An initial concern with such hybrid approaches is that they are liable to render appeal to doxastic factors explanatorily redundant. But the more pressing concern is that such hybrid accounts seem liable to inherit the various problems associated with those other explanatory approaches (see e.g. Nagel 2010a: 286-301, Blome-Tillmann 2013, Dimmock & Huvenes 2014, Dinges forthcoming(a) for some discussion of relevant problems).

way to exploit the doxastic approach to explain our judgments even in the examples sketched above—an appeal to egocentric bias.¹⁹

Empirical research indicates that we have trouble making accurate judgments about those in more naïve positions than ourselves. In particular, a significant body of research has shown that, when making judgments about those more ignorant than ourselves, we tend to mistakenly treat them as though they share our knowledge (see e.g. Nickerson 1999, Birch & Bloom 2004). This tendency is often termed ‘epistemic egocentrism’ or ‘the curse of knowledge’. Nagel does *not* propose that a tendency to share our knowledge might explain our judgments in the cases presented in the previous section, but she does suggest that a similar bias could be in play, since the broader problem of epistemic egocentrism concerns not just our knowledge but also our “beliefs, attitudes and concerns” (2010a: 302).²⁰ Nagel (2010a: 301-6; see also 2010b: 425-6) proposes that our judgments about examples like *Table (3rd Person)* and *Table (Modal Contrast)* might be a reflection of an egocentric tendency to treat others as sharing our concerns about error possibilities. In regard to a case similar to *Table (3rd Person)*, Nagel writes:

Once concerns about the possibility of tricky lighting have been raised for me, I illegitimately evaluate [the subject’s] predicament as if he shared those concerns. Ordinarily, a person who is actively concerned about the lighting would glance up to check it prior to making a judgement about the colour of the table. People who are actively thinking about the influence of lighting conditions on colour judgements can still go ahead and

¹⁹ Nagel is not the only classical invariantist to invoke psychological bias to explain problematic judgments. Williamson 2005 and Gerken 2012 (see also Gerken & Beebe 2016) also defend classical invariantism via appeal to psychological bias. For criticism of Williamson’s approach, see Nagel (2010a: 286-301); for criticism of Gerken, see Stoutenburg (2017).

²⁰ Nagel does not explicitly cite any literature in support of this claim. Although it is widely accepted that we tend to treat others as sharing our beliefs, attitudes and concerns (see e.g. the literature on the false consensus effect (Ross *et al.* 1977; Dawes 1989)), it is less clear that these tendencies will all have the same characteristics as epistemic egocentrism (understood narrowly as a tendency to treat others as sharing our knowledge). In particular, it is less clear whether tendencies to treat others as sharing our ‘beliefs, attitudes and concerns’ will be as robust as our tendency to treat others as sharing our knowledge. (Of potential relevance here: see Birch and Bloom (2004: 257-8; also 256, Box 1) on the contrast between treating others as sharing our knowledge vs. sharing our ignorance.) For the purposes the present paper, I shall just grant to Nagel that the relevant tendencies are equally robust.

make their colour judgements without checking the lighting, but would typically do so only under conditions of compromised or motivated belief formation. (Nagel 2010a: 303)

By reading the description in *Table (3rd Person)*, the possibility that the table is white but illuminated by red lights becomes salient to us. Due to egocentric bias, we mistakenly treat Alice as though she is also considering that possibility. John's knowledge denial to Alice therefore seems true because it seems to us that Alice could only hold the belief that the table is red if she were under the influence of some epistemically problematic factor, such as wishful thinking.²¹

An initial concern. Is it plausible that we would treat Alice as though she is considering the possibility that the table is white but illuminated by red lights even though it is apparent from the case description that she is not considering such error possibilities? An interesting feature of epistemic egocentrism is that it is surprisingly robust. We tend to treat others as sharing our knowledge (at least to some degree) even when it is apparent that they lack this knowledge (see e.g. Fischhoff 1975, Camerer *et al.* 1989, Birch & Bloom 2004). Nagel (2010a: 301-6) proposes that our tendency to treat others as considering the error possibilities that we are considering might exhibit similar robustness—that we exhibit the relevant bias even if it is stipulated that the subject is not considering such possibilities. For the present, let us grant that this response is successful; the issues here will come to the fore in §4-6.

A similar egocentric explanation can be offered in response to examples like *Table (Modal Contrast)*. In regard to that case, the proposal would be that when trying to imagine counterfactual situations in which John is not considering the possibility that the table is

²¹ Nagel (2010b: 425-6) offers an alternative explanation for stakes-based cases similar to *Table (3rd Person)*. Her response mirrors one found in Stanley (2005: 102-4). For criticism of that proposal, see Schaffer (2006: 93-4) & MacFarlane (2014: 186-7). Bach (2005: §V) also offers an alternative error-theoretic treatment of the cases; I consider Bach's response in fn. 37.

white but illuminated by red lights, we still mistakenly treat John as though he is considering that possibility. As a result, it seems to us that even if John were not considering the possibility that the table is white but illuminated by red lights, he would still fail to know that the table is red, since he would only be able to form the belief that it is red if he were under the influence of some epistemically problematic factor, like wishful thinking. Thus, Mona's utterance of 'He doesn't know that the table is red. But if his son hadn't raised the possibility that the table is white but illuminated lights, he would know it's red' seems false because her utterance of the conditional seems false (cf. Nagel 2010b: 426).²²

The proposal also seems able to account for the more general observation associated with our problem cases—namely, that when the possibility that the table is white but illuminated by red lights becomes suitably salient, we are prone to treat both subjects who are and subjects who are not considering that possibility as needing to be able to rule out that possibility in order to know. If, due to egocentric bias, we treat subjects who it is apparent are not considering the error possibilities that we are as though they are considering those possibilities, appeal to doxastic effects can presumably explain why we treat those subjects as also needing to rule out the error possibilities we are considering in order to possess knowledge.²³

²² What about structurally similar cases that concern shifts in practical factors, like stakes, rather than shifts in salient error possibilities (see e.g. Stanley 2005: 3-6 & 106 for relevant cases)? Nagel (2010b) proposes that high stakes subjects exhibit higher levels of 'epistemic anxiety' than low stakes subjects. (A subject's level of epistemic anxiety corresponds (roughly) to the amount of evidence the subject needs to possess in order to be able to naturally form the relevant belief.) To handle stakes-based versions of examples like *Table (3rd Person)* and *Table (Modal Contrast)*, Nagel (2010b: 425-6) suggests that, due to egocentric bias, high stakes subjects may be prone to treat low stakes subjects as though they share their own high levels of epistemic anxiety. The concerns to follow carry over fairly straightforwardly to this kind of egocentric strategy as well (cf. fn. 33). (Nagel (2008: 292) offers a slightly different egocentric bias strategy for handling so-called 'Ignorant High Stakes' cases; the concerns raised below may also pose a problem for this explanation, but for considerations of space, I cannot pursue the issue here.)

²³ Stoutenburg (2017: 2037-9) objects to Nagel's appeal to epistemic egocentrism on the grounds that she has not first explained why we treat subjects who are considering (e.g.) the possibility that the table is white but illuminated by red lights as needing to rule out that possibility in order to know that the table is red. However, Stoutenburg does not mention or consider the doxastic elements of Nagel's proposal that are intended to handle this issue. Roughly, Nagel's suggestion is that, due to the psychological constraints on belief formation, subjects who are considering the possibility of tricky lighting will not be psychologically able to form the belief that the table is red (without the

Interestingly, appeal to egocentric bias would also seem to have the potential for much broader philosophical application. For one thing, it seems that structurally similar appeals to egocentric bias could be put to work defending rival positions in the debates surrounding the metaphysics/semantics of knowledge. For example, *impurist invariantists* (commonly called ‘subject-sensitive invariantists’) claim that knowledge is necessarily connected to the subject’s conversational or practical concerns. On such accounts, the ignorance-inducing effects of considering error possibilities, or focusing on particular practical matters, are limited to those subjects who are in fact considering those error possibilities or focusing on those practical matters. As a result, impurist invariantists also face trouble with examples like *Table (3rd Person)* and *Table (Modal Contrast)* (see e.g. Stanley 2005: 98-9 & 106). But drawing on Nagel’s suggestions, impurist invariantists might propose that our judgments regarding such cases are not in fact a threat to their theories: they are merely the upshot of an egocentric tendency to treat other subjects as though they share our conversational and practical concerns.²⁴

There is also potential for application in other domains. For example, a lot has been made in the literature on epistemic modals (epistemic uses of expressions like ‘might’ and ‘possible’) of so-called ‘eavesdropper cases’. These are examples in which an eavesdropper, who knows that P, overhears a speaker, who does not know that P, make a claim of the form ‘It might be that not-P’. It is oft-reported that it seems true for the eavesdropper to say ‘That’s false’ in reference to the speaker’s ‘might’-claim (see e.g. Hawthorne 2007). Such cases pose a problem for some standard contextualist accounts of epistemic modals on which epistemic uses of ‘It might be that not-P’ express, roughly, that not-P is compatible with what the speaker knows (or with what the speaker and her intended audience knows).

influence of epistemically problematic factors) unless they have evidence sufficient to rule that possibility out (see §1-2 above & esp. Nagel 2011: 13-5).

²⁴ See Dinges forthcoming(b) for a recent development of this kind of proposal. I take the concerns raised in §4-6 to carry over fairly straightforwardly to Dinges’ proposal, but for considerations of space, I cannot engage with his proposal here. Thanks to an anonymous referee for drawing my attention to Dinges’ paper.

But if the eavesdropper (and us) treat the speaker as sharing our knowledge that P, that could explain why the eavesdropper (and us) judge that it seems correct for the eavesdropper to say ‘That’s false’.

Egocentric bias could also be central to explaining some cases of moral disagreement and of disagreement about personal taste that pose trouble for contextualist theories of moral terms like ‘good’ and ‘wrong’ (see e.g. Khoo & Knobe forthcoming) and of taste expressions like ‘fun’ and ‘disgusting’ (see e.g. MacFarlane 2014: ch. 6). If we treat others as sharing our moral codes and personal tastes, then even supposing that terms like ‘good’ and ‘disgusting’ express different properties when used by different those with different moral codes or personal tastes, we may mistakenly treat others as expressing the same properties that we express when using those terms—and erroneously judge that we are in disagreement as a result.²⁵ Given its potentially broad diagnostic potential, then, it seems that a better understanding of egocentric bias would be beneficial.

§4 A Partial Bias

Nagel suggests that our tendency to treat others as considering the error possibilities that we are considering is a facet of epistemic egocentrism (Nagel 2010a: 301-6, 2010b: 425). This is important for Nagel’s project because epistemic egocentrism is a surprisingly robust bias. Remarkably, subjects in the relevant empirical studies continue to be biased even if they are told about epistemic egocentrism and its effects on judgment (Fischhoff 1975), and even if they are given financial incentives to avoid the bias (Camerer *et al.* 1989). But although the

²⁵ The point here is that appeal to egocentric bias appears to show significant *prima facie* explanatory promise. The extent to which appeal to egocentric bias can assist in explaining our judgments seems liable to depend, *inter alia*, on the precise details of the contextualist accounts at issue. Similar remarks apply to the other potential applications of egocentric bias sketched above.

bias is robust, it is not as strong as our previous discussion (and Nagel's own discussion) might have seemed to indicate.

In regard to very young children, Birch and Bloom (2004, 2007) suggest that epistemic egocentrism *is* very powerful. Consider an example familiar from debates in developmental psychology:

Displacement. An experimental participant is told a story about 'Sally' who places her candy in basket A and then leaves the room. In the story, another character then moves Sally's candy to basket B. The participant is asked where Sally will first look for her candy when she comes back into the room. Children under four typically respond that Sally will first look in basket B. (Wimmer & Perner 1983; Baron-Cohen, Leslie, & Frith 1985.)

Birch & Bloom (2004, 2007) (see also Birch & Bernstein 2007) suggest that these results be understood in terms of epistemic egocentrism. Birch & Bloom propose that the child responds that Sally will first look for her candy in basket B because the child treats Sally as sharing his knowledge that the candy is in basket B, and that he does this despite being apprised of the information (from the story) that Sally was out of the room when her candy was moved.²⁶

This clearly resembles the kind of bias that Nagel alleges is present in examples like *Table (3rd Person)* and *Table (Modal Contrast)*. We treat the relevant subject as sharing our concern with the possibility that the table is white but illuminated by red lights, despite being apprised of the information that the subject is not considering such error possibilities.

²⁶ The traditional understanding of these tasks is that they reveal a more fundamental cognitive deficit in very young children than simply epistemic egocentrism. Birch & Bloom (2004, 2007) (also Birch & Bernstein 2007) attempt to push back against that traditional understanding, but the dispute is not important for our purposes. The central point is that epistemic egocentrism does not manifest in such a powerful way in adults.

The problem is that epistemic egocentrism does not manifest in this kind of strong way in adults.

A central feature of epistemic egocentrism as it manifests in adults is that it is “a partial bias” (Birch & Bloom 2004: 258, Box 2; see also Camerer *et al.* 1989). We (adults) do not straightforwardly treat others as though they know what we do, at least not when it is apparent that those others do not share our knowledge. Rather, our knowledge impairs our ability to make accurate judgments about those more ignorant than ourselves. A common theme in the literature on understanding egocentric bias is the idea that our judgments about how another person will act or judge are shaped by our thinking about how we would act or judge in that person’s situation (see e.g. Nickerson 1999; cf. Nagel 2010a: 302). On this picture, epistemic egocentrism with respect to knowledge emerges because we are unable to fully suppress the effects of our own knowledge when making judgments about those more ignorant than ourselves. The result is various partial errors.

Consider the displacement task. Adults (and older children) do not judge that Sally will first look in basket B. We thus do not treat Sally as though she straightforwardly shares our knowledge of her candy’s location when making a judgment about how Sally will act. But research does suggest that our own knowledge that the candy is in basket B might lead us to commit various partial errors.

For example, one potential partial effect is that our own knowledge that the candy is in basket B will lead us to overestimate the *likelihood* that Sally will first look in basket B. In this regard, Birch and Bloom (2007) conducted a study that indicates that (adult) subjects who are told that Sally’s candy is removed from basket A and then placed in basket B will judge that it is more likely that Sally will first look in basket B than will those subjects who

are told merely that the candy has been removed from basket A and then returned to a basket (but not told which basket).²⁷ Such results indicate that the more knowledgeable subjects are biased in their judgments about the likelihood of where Sally will first look for her candy. But any bias in their judgments about how Sally will act is decidedly partial: Birch & Bloom's study suggests that the more knowledgeable subjects will still assign a low probability to the claim that Sally will first look in basket B, even if it is higher than the probability assigned by the more ignorant subjects.

A couple of prominent examples serves to further demonstrate the partial nature of epistemic egocentrism. An influential study by Baron and Hershey (1988), one Nagel (2010a: 302-3) emphasises when introducing her own proposal, comprised an investigation into the effects of outcome knowledge on our evaluation of various medical and monetary decisions. As part of the study, subjects were given the following medical case:

A 55-year old man had a heart condition. He had to stop working because of chest pain. He enjoyed his work and did not want to stop. His pain also interfered with other things, such as travel and recreation. A type of bypass operation would relieve his pain and increase his life expectancy from age 65 to age 70. However, 8% of the people who have this operation die from the operation itself.

His physician decided to go ahead with the operation. The operation succeeded. Evaluate the physician's decision to go ahead with the operation. (Baron & Hershey 1988: 57)

Participants were also supplied with a case that was the same except that it involved a negative outcome—in that case, it was stated that the operation was unsuccessful and the

²⁷ Birch and Bloom's (2007) study was conducted on a more complicated example than the one discussed in the main text. Their example involved several baskets of different shapes and colours, and the baskets themselves were also moved around while the Sally-character was out of the room. The results showed statistically significant bias in the judgments of subjects who knew where the object (a violin) was placed: those participants who knew which basket the violin was in judged it more likely that the Sally-character would first look in that basket than did those participants ignorant of the violin's location. (Though interestingly Birch and Bloom found no statistically significant bias in versions of the example where it was especially implausible that Sally would first look in the basket where the violin in fact was. To the extent that it is especially implausible that a typical subject would be considering the possibility that the table is white but illuminated by red lights, this may be a source of additional concern for Nagel's proposal. But I set it aside.)

man died.²⁸ The participants were told that the man's physician had no further information on which to ground her decision to go ahead with the operation other than what is given in the first paragraph of the quoted text, and were asked to evaluate the physician's decision, "the decision itself, the quality of the thinking that went into it", on the following scale: 3—clearly correct, and the opposite decision would be inexcusable; 2—correct, all-things-considered; 1—correct, but the opposite would be reasonable too; 0—the decision and its opposite are equally good; -1—incorrect, but not inexcusable; -2—incorrect, all-things-considered; -3—incorrect and inexcusable.

If participants had straightforwardly treated the physician as though she shared their knowledge of the actual outcome when making her decision, we would presumably expect a preponderance of 3s in the case where the operation succeeded, and -3s in the case where the man died. The reason is simple: if (e.g.) the physician knew the man was going to die, then her decision to undertake the operation would presumably be incorrect and inexcusable. But this is not what happened. Instead, what was found was a tendency to evaluate the physician's decision as slightly higher up the aforementioned scale in the case where the operation succeeded than in the case where the man died. (Baron & Hershey calculated the mean decision evaluation in the positive outcome (operation succeeded) case and then subtracted it from the mean decision evaluation in the negative outcome (patient died) case. They report that, across a range of similar positive/negative-outcome case pairs, the mean difference was 0.7—i.e. a positive outcome as opposed to a negative outcome produced, on average, an increase of slightly more than half an increment on the ranking scale described above.) This suggests that our evaluation of decisions is indeed biased by our knowledge of the relevant outcome, but that once again, the effect is partial. We do not straightforwardly treat the physician as though she shares our knowledge of the outcome.

²⁸ The cases were presented along with several others in a within-subjects design; the positive outcome case was spread far apart in the presentation from its corresponding negative outcome case to reduce reliance on memory.

Consider a final example. As part of a well-known study by Baruch Fischhoff (1975), a group of participants ('before-subjects') were given some information concerning a particular event (e.g. a conflict between British and Gurka forces in the 1800s), and asked to estimate the likelihood of each of a list of four possible outcomes, given that information. Another group of participants were given the same information, but were also told that a particular outcome in fact came to pass (e.g. that the Gurka forces triumphed). This second group of participants ('after-subjects') were asked to judge the likelihood of the four possible outcomes 'as they would have, had they not known what happened'. After-subjects gave estimates of how likely they would have judged the reported outcome to be that were significantly higher than the estimates produced for that outcome by before-subjects. Nevertheless, when asked to judge 'as they would have, had they not known what happened', after-subjects did not treat their counterfactually-imagined more ignorant selves as though they knew the relevant outcome. If they had, they would presumably have issued probabilities close to 100% for the outcome that they were told came to pass.²⁹ Instead, after-subjects' estimates of how likely they would have judged that outcome to be were merely inflated. (Across a range of historical events, Fischhoff reports that the mean increase in probability estimate for the reported outcome among after-subjects vs. before-subjects was 9.2%.)³⁰

The upshot of such studies is that although we are often 'cursed' by our own knowledge when making judgments about others, we do not straightforwardly treat others as though they know what we know, at least not in situations where it is apparent that they do not possess that knowledge. The claim that we treat such subjects as 'sharing' our knowledge is

²⁹ Participants were asked to express their probability estimates as percentages.

³⁰ Note that unlike the Baron & Hershey (1988) study, Fischhoff's study employed a between-subjects design (as did Birch & Bloom 2007).

thus a rather misleading oversimplification (cf. Birch & Bloom 2004: 258, Box 2). What really happens is that our own privileged knowledge leads us to make various partial judgmental errors—errors in judgment about how others will judge and act, or about the appropriateness of their so doing, that are plausibly the result of a failure to fully suppress our own knowledge. A more accurate characterisation of the phenomenon is thus that we treat those who it is apparent are more ignorant than ourselves as though they share our knowledge *to some degree*.³¹ Or drawing on the familiar language of ‘epistemic position’ (DeRose 1995), it seems that our own knowledge leads us to treat more ignorant others as though they occupy a *better epistemic position* (one more like our own) than they really do. Consider the displacement task. We do not straightforwardly treat Sally as sharing our knowledge that her candy is in basket B. But insofar as we overestimate the likelihood that Sally will first look in basket B, we do seem to be treating Sally as though her epistemic position with respect to the proposition that her candy is in basket B is better than it really is.

§5 Partial Bias and Salient Error Possibilities

Let’s turn now to consider the other facet of epistemic egocentrism (broadly understood) that Nagel proposes accounts for our judgments in the problem cases encountered in §2—that we will treat other subjects as though they are considering the error possibilities that we are. The preceding reflections on how epistemic egocentrism manifests with respect to knowledge suggest that it is implausible that we will straightforwardly treat others as

³¹ In regard to Baron & Hershey’s (1988) investigation into our evaluation of medical and monetary decisions, Nagel herself writes that “the subjects began to misrepresent the decision-makers egocentrically as though they did have *some degree of foreknowledge*” (Nagel 2010a: 303; emphasis added). This passage suggests that Nagel recognises the partial nature of epistemic egocentrism as it manifests with respect to knowledge. It is thus somewhat surprising that she does not acknowledge that the same is likely to be the case with respect to her own proposed bias.

though they are considering the same unlikely error possibilities that we are considering, at least not when it is apparent that they are not considering those possibilities.³²

Consider Alice in *Table (3rd Person)*. It is relatively uncommon for a subject to be considering unlikely error possibilities, and there is no apparent reason to think that Alice would be inclined to do so. It is also explicitly stipulated in the case description that Alice is not considering error possibilities (ways she might be mistaken), such as the possibility that the table is white but illuminated by red lights. The preceding reflections on epistemic egocentrism thus suggest that we will not treat Alice as though she is straightforwardly considering that error possibility. (At least not so long as we are over 4 years old.) Instead, we will exhibit (at most) partial bias. Mirroring how epistemic egocentrism manifests with respect to knowledge, we might expect that we will treat Alice as though she shares our concerns with error possibilities *to some degree*. (We will explore how exactly this partial bias might manifest shortly.)

Similar remarks apply to *Table (Modal Contrast)*. In regard to this case, the Fischhoff (1975) study provides a useful comparison. Just as we do not treat our counterfactually-imagined selves as straightforwardly knowing the outcome when asked to judge ‘as we would have, had we not known the outcome’, so it seems implausible to allege that when asked to consider Mona’s utterance of ‘If his son had not mentioned the possibility that the table is

³² In the example that Nagel (2010a: 287) focuses on, it is not stipulated (or otherwise apparent) that the subject is not considering the possibility that the table is white but illuminated by red lights. In regard to *this* kind of case, it may be plausible to suggest that we tend to straightforwardly treat the subject as also considering the error possibility that we are considering (cf. Nickerson 1999 on how epistemic egocentrism manifests with respect to knowledge). Indeed, Alexander, Gonnerman & Waterman (2014) conducted empirical research that lends some support to the claim that we do engage in straightforward projection of salient error possibilities in such cases (see Nagel & Smith (2017: §5) for some relevant discussion). But obviously the appeal to egocentric bias has very limited application if it can only be used to address cases where the relevant differences between us and the subject are not stipulated to be present (or are not otherwise apparent). (Most immediately, the proposal could not be used to address *Table (3rd Person)* or *Table (Modal Contrast)*.) And note also that Nagel (2010b: 425-6) explicitly seeks to extend the proposal to cases where the relevant differences in conversational or practical concerns *are* stipulated to be present.

white but illuminated by red lights, [John] would know that it's red', we will straightforwardly treat 'counterfactual-John' as though he is considering that error possibility. Once again, the most it seems plausible to expect is partial bias: that we will treat counterfactual-John as though he shares our concerns with error possibilities to some degree.

At this point, those sympathetic to the appeal to egocentric bias might insist that the original proposal was merely an oversimplification—one, perhaps not coincidentally, that is often found in the relevant psychological literature—and that partial bias effects are sufficient to explain our judgments. In the remaining discussion, I shall argue on various grounds that this is not the case, and that Nagel's proposed defence of the doxastic approach is consequently unsuccessful.

On initial inspection, it might seem that partial egocentric effects *are* suited to explaining our judgments. Consider *Table (3rd Person)*. As noted above, our reflections on epistemic egocentrism suggest that we will not straightforwardly treat Alice as considering the possibility that the table is white but illuminated by red lights. But we may nevertheless treat Alice as though she has some greater degree of concern with, or awareness of, error possibilities than she really does. Given the proposals put forward in §1, this seems liable to have an important impact on our assessment of Alice's doxastic condition.

For example, recall *lowers credence*, the view that the consideration of error possibilities tends to result in lower levels of subjective confidence (§1). Given *lowers credence*, if we treat Alice as though she has greater awareness of error possibilities than she really does, we may judge that it would be natural for her to exhibit a lower credence in the proposition that the table is red than would in fact be natural for her to exhibit. As a result, we may be

prone to mistakenly treat her as though she requires more evidence than she really does in order to naturally maintain a credence above the threshold required for belief. Or consider *removes psychological conviction* (§1), the view that the consideration of error possibilities impacts how much evidence a subject requires in order to take the question at issue ‘to be settled’. Given *removes psychological conviction*, if we treat Alice as though she has a greater awareness of error possibilities than she really does, we may be prone to mistakenly treat Alice as requiring more evidence than she really does in order to naturally take it to be settled (and so believe) that the table is red.

More broadly, then, it seems that an egocentric tendency to treat Alice as though she shares our concern about error possibilities to some degree may still lead us to treat Alice as requiring more evidence (or a stronger epistemic position) than she really does in order to naturally form the belief that the table is red. Perhaps this could explain our judgments regarding *Table (3rd Person)*. If we overestimate the evidence that Alice requires in order to naturally form the belief that the table is red, we may mistakenly judge that Alice would not naturally form the belief that the table is red on the evidence she possesses. As a result, it may seem to us that Alice would only form that belief if she were under the influence of some epistemically problematic factor, like wishful thinking. This could explain why John’s utterance of ‘[Alice] doesn’t know that it’s red either’ seems true. Similar remarks apply to *Table (Modal Contrast)*.³³

³³ Note that it is possible to put forward an egocentric proposal that focuses directly on sharing elements of our doxastic condition, rather than on sharing our consideration of error possibilities. For example, it could be proposed that, due to egocentric bias, we directly treat others as though they require the level of evidence that we require in order to take it to be settled that the table is red (cf. Nagel 2010b: 420n). However, so long as it is apparent from the case description that (e.g.) the subject does not require the evidence that we require in order to take it to be settled that the table is red, such alternative proposals will fail for similar reasons as those to be outlined in the next section. The relevant doxastic differences between us and the subject may already be suitably apparent due to the stipulation in the case description that the subject is not considering the error possibilities that we are considering, and also due to it plausibly being common-knowledge that subjects do not typically need to check the lighting before forming colour beliefs. But it does not seem to reverse our judgments if we also make the relevant doxastic differences more explicit. For example, in regard to *Table (3rd Person)*, even if we stipulate in the case description that Alice takes it to be settled that the table is red in a typical automatic way, it still seems plausible that, once the possibility that the table is white but illuminated by red lights has been made

§6 Undermining the Proposal

There are two broad problems with the kind of appeal to partial bias effects sketched at the close of the previous section. The first problem is that there is plausibly an egocentric bias working in the opposing direction to the one Nagel postulates, and that may well cancel out the bias effect that Nagel is proposing.

It is a standard feature of examples like our various '*Table*' cases to stipulate that the relevant proposition is indeed true—e.g. to stipulate that the table is indeed red (as was done in §1). But we are being asked to make judgments about subjects who have not been supplied with this assurance. As our earlier review of the literature on epistemic egocentrism indicated, being explicitly assured of a particular outcome—e.g. that the table is red, that the patient died, that the Gurkas triumphed in their conflict with the British—is itself plausibly going to lead to partial egocentric bias when making judgments about subjects who lack that assurance. In the general terms introduced above, the relevant research would seem to indicate that our own 'God's-eye-view' will lead us to treat the subjects in the various *Table* cases as though their epistemic position with respect to the proposition that the table is red is better (i.e. more like ours) than it really is.

For example, consider *Table (3rd Person)*. When being asked to evaluate a subject like Alice, we are being asked to envisage a subject who *both* lacks our God's-eye-view assurance that the table is indeed red *and* is not considering the error possibilities that we are considering. The preceding reflections suggest that this situation will result in egocentric bias effects acting in two opposing directions. The first egocentric bias effect—

suitably salient, we will judge that John's utterance of 'She has the same evidence as me. She doesn't know either' seems true. Similar remarks apply to *Table (Modal Contrast)*.

tied to our Gods-eye-view—should, broadly speaking, lead us to treat Alice as though her epistemic position is better than it really is. The second effect—tied to our consideration of the possibility that the table is white but illuminated by red lights—should, broadly speaking, lead us to overestimate the strength of epistemic position required for Alice to naturally form the belief that the table is red. The concern is that the first egocentric effect may simply cancel the second one out.

Is a response to this concern available? Nagel does address a similar concern, crediting it to John MacFarlane (Nagel’s text has been amended to fit the present example):

Objection: to the extent that we overshare our privileged information, why wouldn’t we also project onto [the subject] our knowledge that [the table is red]? When we are evaluating [the subject’s] state of mind with respect to this very proposition, we need to represent her as making some transition from her resources to this target. We may naturally misrepresent these resources, but it is unlikely that we would do so by including among them the stipulation about the target proposition that we have been given in our “God’s eye” view of the case: a stipulation of the truth of the proposition cannot be represented as included in her resources for making a judgment on that very proposition. (Nagel 2010b: 425n)

One might question whether we really do represent Alice as making some ‘transition’ from her resources (presumably something like her evidence) to her belief that the table is red when judging whether John’s utterance of ‘She doesn’t know the table is red’ is true. But even setting that issue aside, the more important problem with this response is that it does nothing to challenge the claim that a partial egocentric effect associated with our God’s-eye-view will be present.

Consider the Fischhoff (1975) study mentioned above. In that study, ‘after-subjects’ were given a God’s-eye-view assurance of the actual outcome, and asked to judge ‘as they would have, had they not known the outcome’. Paralleling Nagel’s contention in the quoted passage, it is presumably unlikely that these after-subjects will treat their counterfactually-

imagined more ignorant selves as though the stipulation given in their God's-eye-view forms part of their 'resources'. And indeed that is just what Fischhoff found. As discussed above, his results did not indicate that after-subjects treat their counterfactually-imagined selves as straightforwardly sharing their God's-eye-view of the outcome. But Fischhoff's results nevertheless indicated that the after-subjects' God's-eye-view gave rise to partial egocentric bias effects. In the present terms, after-subjects' responses indicated that they were judging as though their counterfactually-imagined more ignorant selves possessed resources (evidence) that were better than they really were.

In a similar vein, then, although it seems reasonable for Nagel to allege that we will not treat Alice as straightforwardly sharing our God's-eye-view assurance that the table is red, research on epistemic egocentrism nevertheless indicates that our possession of that assurance will give rise to partial egocentric bias. In the present terms, to a tendency to treat Alice as though her resources (evidence) are better than they really are. Nagel's response in the quoted passage therefore does nothing to undermine the claim that a partial bias associated with our God's-eye-view will be present. Given that Nagel's proposed egocentric bias effect centred around the consideration of error possibilities is—as argued above—*also* at most merely partial, the concern that the two partial effects may simply cancel one another out remains in force.³⁴

³⁴ Nagel (2010a: 301) suggests that egocentric bias effects are stronger when the fact that the subject differs in the relevant respect (e.g. considering the possibility of tricky lighting) is not in focus. Could this form the basis for alleging that the egocentric effect associated with our consideration of error possibilities is likely to be stronger than the egocentric effect associated with our God's-eye-view assurance that the table is red? The outlook for such a response is poor. Most immediately, it is far from clear that the fact that the relevant subjects in our 'Table' cases are not considering the error possibilities that we are considering is less in focus than is the fact that those subjects do not share our God's-eye-view assurance that the table is red. In this regard, note (e.g.) that the fact that Alice in *Table (3rd Person)* is not considering error possibilities is explicitly stated in the case description, whereas the fact that Alice does not share our God's-eye-view is not. It also seems especially difficult to pursue this strategy in regard to examples like *Table (Modal Contrast)*. In these kinds of examples, we are being asked to make a judgment about conditionals like 'If those error possibilities hadn't been mentioned, the subject would know' or 'If the subject had not been considering error possibilities, he would know'. When making judgments about these sorts of conditionals, the fact that the subject in the counterfactual situation is not going to be considering the error possibilities that we are seems to be *very* much in focus; it seems hard to argue that the fact that the subject in the counterfactual situation lacks our God's-eye-view assurance is significantly *more* in focus when making judgments about such conditionals.

Let's now turn to look at a second problem: that even if Nagel's proposed egocentric effect is not cancelled out, partial bias effects are in any case insufficient to explain our judgments. An initial worry here concerns just how 'partial' the relevant biases appear to be. Recall the extent of the egocentric bias in the studies from the 'curse of knowledge' literature. In the medical decision case (Baron & Hershey 1988), the subjects told that the patient died, or that the operation succeeded, don't treat the physician as though she knew what they know—they do not evaluate her decision as a -3 (incorrect and inexcusable) or a +3 (correct and the opposite would be inexcusable) on the scale introduced earlier. The effects of being supplied different outcome information, although significant, were much more modest. On average, the difference between knowledge of a positive vs. a negative outcome corresponded to a 0.7 difference on the decision evaluation scale—i.e. a difference of slightly more than $\frac{1}{2}$ a point, far less than the 6 point difference between +3 and -3. Fischhoff's (1975) study on the effects of outcome knowledge on people's estimates of how 'they would have judged, had they not known the outcome', and Birch & Bloom's (2007) study on adults' judgments in displacement tasks, revealed similar egocentric effects: clearly significant, but nothing like the effects that would be expected if we were treating the more ignorant subjects as though they straightforwardly shared our knowledge.

These results raise the concern that any egocentric tendency to treat others as 'sharing' our concerns about error possibilities is likely to be similarly modest. In terms of the effect on our evaluation of Alice's doxastic condition, such results suggest that even if we are prone to overestimate how much evidence Alice requires to naturally form the belief that the table is red, our overestimate is likely to be relatively modest. However, the evidence that Alice in fact possesses—the evidence acquired by virtue of her clear view of the table from a few feet away in apparently normal lighting conditions—is plausibly more than sufficient for a

subject (like Alice) who is not considering error possibilities to naturally form the belief that the table is red. Weaker evidence, such as the evidence acquired from perception at a greater distance, or from testimony from a generally reliable source, would presumably be sufficient.³⁵ This raises the concern that if our overestimate of the evidence that Alice requires in order to naturally form the belief that the table is red is relatively modest, it may well still strike us that the evidence Alice in fact possesses is sufficient for her to naturally form the belief that the table is red. We should therefore remain sceptical that appeal to partial egocentric bias is sufficient to explain our judgments regarding examples like *Table (3rd Person)*. (A similar concern arises in regard to examples like *Table (Modal Contrast)*.)

There is however a more pressing problem. Consider John in *Table (3rd Person)*, a subject who is straightforwardly considering the possibility that the table is white but illuminated by red lights. Plausibly, the minimum additional evidence that strikes us as required for John to be truly said to ‘know’ that the table is red is the minimum evidence sufficient to rule out that error possibility—i.e. something like the evidence acquired by looking up and checking the lighting (see §2). To account for this, advocates of the doxastic approach should presumably allege that the minimum additional evidence that we take John to require to naturally form the belief that the table is red is the minimum evidence sufficient to rule out the possibility that the table is white but illuminated by red lights (cf. Nagel 2010a: 303, 2011: 13-5).

As argued above, an upshot of the partial nature of egocentric bias is that we won’t treat Alice as just like John—we won’t treat her as though she is straightforwardly considering the possibility that the table is white but illuminated by red lights. In terms of the effect on our assessment of Alice’s doxastic condition, this suggests that whatever additional

³⁵ See Nagel (2010b) for some relevant discussion concerning the evidence required for belief formation among subjects in different cognitive conditions.

evidence (if any) we take Alice to require to naturally form the belief that the table is red is going to be weaker than whatever additional evidence we take John to require to naturally form that belief. Assuming the minimum additional evidence that we take John to require is the minimum evidence sufficient to rule out the possibility that the table is white but illuminated by red lights, it should therefore seem to us that Alice is able to naturally form the belief that the table is red on weaker evidence than is required to rule out that possibility.

This raises the concern (*inter alia*) that even if appeal to egocentric bias can explain our judgments regarding the original *Table (3rd Person)* case, it is not going to be able to explain judgments regarding a range of cases almost identical to *Table (3rd Person)*. In particular, appeal to egocentric bias seems ill-suited to explain our judgments regarding a range of cases that are the same except that both John and Alice have better evidence that the table is red—but in which they still do not have the kind of evidence intuitively required to rule out that the table is white but illuminated by red lights. In such cases, John's son will raise the possibility that the table is white but illuminated by red lights, and it is presumably still going to seem true both for John to utter 'I don't know that the table is red' and for him to go on to claim 'Alice has the same evidence as me. She doesn't know either' (cf. Bach 2005: §V). However, appeal to egocentric bias is not going to be sufficient to explain why we would take Alice to lack knowledge in all of these cases, since, across a range of these cases, Alice presumably ought to strike us as having evidence sufficient to naturally form the belief that the table is red. (Once again, parallel concerns arise regarding *Table (Modal Contrast)*.)

Our reflections on the partial nature of epistemic egocentrism thus bring out a number of pressing concerns. The first main concern is that there will be an egocentric bias acting in the opposing direction to the bias that Nagel postulates, and that may well cancel that bias

out. A further set of problems centres on the explanatory adequacy of partial bias effects. There is an initial worry: that due to the decidedly partial nature of egocentric bias, it is far from guaranteed that Nagel's proposed egocentric effect is going to be sufficient to explain our judgments in the original cases. But there is also a more pressing problem: that even if the effect is suited to explain our judgments in the original cases, there are likely to be structurally similar cases that appeal to egocentric bias cannot explain.³⁶

§7. Closing remarks

The preceding reflections on egocentric bias suggest that examples involving subjects in different conversational situations from the speaker remain problematic for doxastic approaches. Although it has not been argued for explicitly here, it also seems likely that structurally similar concerns will undermine attempts to appeal to egocentric bias to explain similar problematic cases involving subjects in different practical situations from the speaker (see fn. 22). In the absence of a satisfying alternative explanation of these various cases, the doxastic approach looks to be in trouble.³⁷ Given the problems facing other

³⁶ Advocates of the doxastic approach are also apparently left unable to respect the more general observation made at the close of §2: that once an error possibility has become suitably salient, we are prone to judge as though both subjects who are *and* subjects who are not considering that error possibility must be able to rule it out in order for them to be truly said to 'know'.

Nagel & Smith (2017: §5) raise another potential concern with Nagel's appeal to egocentric bias: that if the proposal is right, we should judge that a subject like Alice lacks both knowledge and justified belief, but we are only tempted to judge that such a subject lacks knowledge. I cannot assess this concern here, but note that it seems to rest *inter alia* on the assumption that the doxastic condition on knowledge is the same as the state picked out by our ordinary use of 'belief'; it is possible to pursue a doxastic approach to explaining our judgments about 'know' without endorsing that assumption (see fn. 9). See also Pynn (2014: 129-30) for some brief criticism of Nagel's proposal.

³⁷ As noted earlier (fn. 21), Nagel does offer an alternative explanation for some cases similar to *Table (3rd Person)*, but that approach has met with criticism (and is limited in scope). Bach (2005: §V) offers a different response to cases like *Table (3rd Person)* and *Table (Modal Contrast)*. He suggests (roughly) that we treat the amount of evidence that we require in order to form a *belief* as the amount that other people require in order to *know*—and do so even when (as often happens) the amount required for others to know varies very significantly from how much we require in order to form the relevant belief. For example, consider *Table (3rd Person)*. Suppose that we require evidence sufficient to rule out that the table is white but illuminated by red lights in order to believe that it is red. Bach's proposal is that we will then (mistakenly) treat Alice as requiring that much evidence in order to know that it is red. But why would we do this? As far as I can see, Bach provides no adequate answer to this question. (Interestingly, one might try to appeal to egocentric bias to explain why we treat others as requiring the evidence that we require, and thereby attempt to fill the explanatory hole in Bach's account. However, structurally similar concerns about partial bias will plausibly undermine such a strategy.)

classical invariantist explanations of the shifts in our judgments about ‘know’, this is surely unwelcome news for those seeking to preserve a more conservative epistemological outlook.

Our discussion of egocentric bias also holds some more general philosophical lessons. The partial nature of egocentric bias suggests that its potential for broader diagnostic application may be more limited than indicated in our earlier discussion (§3). For example, consider the eavesdropper cases mentioned above, which have proven central to the literature on epistemic modals. Recall that in those cases, we judge that an eavesdropper who knows that P speaks truly when she utters ‘That’s false’ in reference to a more ignorant speaker’s utterance of ‘It might be that not-P’. If we straightforwardly treated the speaker (and her intended audience) as sharing our knowledge that P, that could explain why we judge the eavesdropper’s claim to be true, even supposing that what the speaker’s ‘might’ claim expresses is that not-P is compatible with what the speaker (and her intended audience) knows. But our discussion indicates that so long as it is clear from the case description that the speaker and her intended audience do not share our knowledge that P, we will not treat them as though they do share that knowledge. (Instead, we will at most treat the more ignorant speaker (and her intended audience) as though they share our knowledge to some degree, or exhibit a better epistemic position with respect to P than they really do.) This suggests that there will be a broad range of eavesdropper cases that appeal to egocentric bias cannot assist in explaining.

Nevertheless, our reflections should not be taken to imply that egocentric bias is entirely unimportant to philosophical practice and theorising. For one thing, it seems like a kind of biasing effect that philosophers should keep in mind when constructing hypothetical cases and testing them on ordinary speakers. And appeal to egocentric bias may well still have

crucial roles to play in explaining our judgments regarding hypothetical cases—we just need to be sure to take account of the partial ways in which the bias manifests when making such explanatory moves.*

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). London: Bloomsbury.
- Bach, K. (2005). The Emperor's New 'Knows'. In G. Preyer & G. Peter (Eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth* (pp. 51-89). Oxford: Oxford University Press.
- Baron, J. & Hershey, J.C. (1988). Outcome Bias in Decision Evaluation. *Journal of Personality and Social Psychology*, 54, 569-79.
- Baron-Cohen, S., Leslie, A.M. & Frith, U. (1985). Does the autistic child have a "theory of mind"? *Cognition*, 21, 37-46.
- Birch, S. & Bernstein, D. (2007). What can children tell us about hindsight bias: a fundamental constraint on perspective-taking? *Social Cognition*, 25, 98-113.
- Birch, S. & Bloom, P. (2004). Understanding Children's and Adults' Limitations in Mental State Reasoning. *Trends in Cognitive Sciences*, 8, 255-260.
- Birch, S. & Bloom, P. (2007). The curse of knowledge in reasoning about false beliefs. *Psychological Science*, 18, 382-6.
- Blome-Tillmann, M. (2009). Contextualism, Subject-Sensitive Invariantism, and the Interaction of 'Knowledge'-Ascriptions with Modal and Temporal Operators. *Philosophy and Phenomenological Research*, 79, 315-331.
- Blome-Tillmann, M. (2013). Knowledge and Implicatures. *Synthese*, 190, 4293-4319.
- Blome-Tillmann, M. (2014). *Knowledge and Presupposition*. Oxford: Oxford University Press.
- Buckwalter, W. (2014). The Mystery of Stakes and Error in Ascriber Intuitions. In J. Beebe (ed.), *Advances in Experimental Epistemology* (pp. 145-74). London: Bloomsbury.

* Thanks to Jessica Brown, Yuri Cath, Torfinn Huvenes, Michael Lynch, Daniele Sgaravatti, and an anonymous referee for helpful comments and discussion.

- Buckwalter, W. & Schaffer, J. (2014). Knowledge, Stakes, and Mistakes. *Noûs*, 49, 201-34.
- Camerer, C., Loewenstein, G., and Weber, M. (1989). The curse of knowledge in economic settings. *Journal of Political Economy*, 97, 1232-1254.
- Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives*, 13, 57-89.
- Dawes, R. (1989). Statistical Criteria for Establishing a Truly False Consensus Effect. *Journal of Experimental Social Psychology*, 25, 1-17.
- DeRose, K. (1995). Solving the Skeptical Problem. *The Philosophical Review*, 104, 1-51.
- DeRose, K. (2004). The Problem with Subject Sensitive Invariantism. *Philosophy and Phenomenological Research*, 68, 346-350
- DeRose, K. (2009). *The Case for Contextualism, vol. 1*. Oxford: Clarendon Press.
- Dimmock, P. & Huvenes, T. Knowledge, Conservatism, and Pragmatics. *Synthese*, 191, 3239-69.
- Dinges, A. (forthcoming a). Knowledge, Intuition, and Implicature. *Synthese*.
- Dinges, A. (forthcoming b). Anti-intellectualism, Egocentrism, and Bank Case Intuitions. *Philosophical Studies*.
- Fantl, J. & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford: Oxford University Press.
- Fischhoff, B. (1975). Hindsight is not equal to foresight: the effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, 1, 288–299.
- Gerken, M. (2012). Epistemic Focal Bias. *Australasian Journal of Philosophy*, 91, 41-61.
- Gerken, M. & Beebe, J. (2016). Knowledge in and out of Contrast. *Noûs*, 50, 133-64.
- Gopnik, A. & Aslington, J.W. (1988). Children's understanding of representational change and its relation to the understanding of false belief and the appearance-reality distinction. *Child Development*, 59, 26–37.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Clarendon Press.
- Hawthorne, J. (2007). Eavesdroppers and Epistemic Modals. *Philosophical Issues*, 17, 92-101.
- Ichikawa, J.J. (2017). *Contextualising Knowledge: Epistemology and Semantics*. Oxford: Oxford University Press.
- Kelley, H. (1972). Attribution in Social Interaction. In Jones, E., Kanouse, D., Kelley, H., Nisbett, R., Valins, S., & Weiner, B. (Eds.), *Attribution: Perceiving the Causes of Behavior* (pp. 1-26). Morristown, NJ: General Learning Press.

- Khoo, J. & Knobe, J. (forthcoming). Moral disagreement and Moral Semantics. *Noûs*.
- MacFarlane, J. (2005). The Assessment Sensitivity of Knowledge Attributions. In T.S. Gendler & J. Hawthorne (eds.), *Oxford Studies in Epistemology: Volume 1* (pp. 197-233). Oxford: Oxford University Press.
- MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and its Applications*. Oxford: Clarendon Press.
- Nagel, J. (2008). Knowledge Ascriptions and the Psychological Consequences of Changing Stakes. *Australasian Journal of Philosophy*, 86, 279-294.
- Nagel, J. (2010a). Knowledge Ascriptions and the Psychological Consequences of Thinking about Error. *The Philosophical Quarterly*, 60, 286-306.
- Nagel, J. (2010b). Epistemic Anxiety and Adaptive Invariantism. *Philosophical Perspectives*, 24, 407-435.
- Nagel, J. (2011). The Psychological Basis of the Harman-Vogel Paradox. *Philosophers' Imprint*, 11, 1-28.
- 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). Abingdon: Routledge.
- Nickerson, R. (1999). How We Know – and Sometimes Misjudge – What Other People Know: Imputing One's Own Knowledge to Others. *Psychological Bulletin*, 125, 737-759.
- Neta, R. (2007). Anti-Intellectualism and the Knowledge-Action Principle. *Philosophy and Phenomenological Research*, 75, 180-87.
- Pynn, G. (2014). Unassertibility and the Appearance of Ignorance. *Episteme*, 11, 125-143.
- Ross, L., Greene, D. & House, P. (1977). The “false consensus effect”: an egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology*, 13, 279–30.
- Schaffer, J. & Knobe, J. (2012). Contrastive Knowledge Surveyed. *Noûs*, 46, 675-708.
- Shin, J. (2014). Time Constraints and Pragmatic Encroachment on Knowledge. *Episteme*, 11, 157-180.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford: Clarendon Press.
- Stoutenberg, G. (2017). Strict Moderate Invariantism and Knowledge Denials. *Philosophical Studies*, 174, 2029-44.
- Vogel, J. (1990). Are there Counter-examples to the Closure Principle? In M. Roth and G. Ross (eds.), *Doubting: Contemporary Perspectives on Skepticism* (pp. 13-28). Dordrecht: Kluwer.

Weatherson, B. (2005). Can we do without pragmatic encroachment? *Philosophical Perspectives*, 19, 417-43.

Weatherson, B. (2017). Interest-Relative Invariantism. In J.J. Ichikawa (Ed.), *The Routledge Handbook of Epistemic Contextualism* (pp. 240-54). Abingdon: Routledge.

Williamson, T. (2005). Contextualism, Subject-sensitive Invariantism and Knowledge of Knowledge. *The Philosophical Quarterly*, 55, 213-35.

Wimmer, H. & Perner, J. (1983). Beliefs about beliefs: representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition*, 13, 103-28.