

COMPARING CONTEXTUALISM AND INVARIANTISM ON THE CORRECTNESS OF CONTEXTUALIST INTUITIONS

Jessica BROWN
University of Bristol

Summary

Contextualism is motivated by cases in which the intuitive correctness of a range of phenomena, including knowledge attributions, assertions and reasoning, depends on the attributor's context. Contextualists offer a charitable understanding of these intuitions, interpreting them as reflecting the truth value of the knowledge attributions and the appropriateness of the relevant assertions and reasoning. Here, I investigate a range of different invariantist accounts and examine the extent to which they too can offer a charitable account of the contextualist data

0. *Introduction*

Contextualism is motivated by cases in which the intuitive correctness of a range of phenomena, including knowledge attributions, assertions and reasoning, depends on the attributor's context. Contextualists take such cases to show that the truth conditions of knowledge attributions depend on the attributor's context. For instance, they suggest that, in the high context of such cases, the denial of knowledge to the subject seems correct because it is literally true. Invariantists deny that the truth conditions of knowledge attributions depend on the attributor's context. So, it may seem that invariantists need to treat our intuitions about contextualist cases as misguided. For instance, if invariantists hold that, in the high context of the contextualist cases, the denial of knowledge is literally false then, on their view, our intuition that the denial is correct seems misguided.

Prima facie, if invariantism treats our intuitions about contextualist cases as misguided, this may seem to count against the view. DeRose argues that it is objectionable for an account of 'knows' to treat as

misguided our intuitions about simple applications of the term, suggesting that ‘what speakers know best about the piece of language at issue is how and when to make simple positive and negative applications of the term in question’ (DeRose *forthcoming*, 24). Some suggest this *prima facie* consideration may be counterbalanced by considering our intuitions about more complex uses of ‘knows’. They suggest that invariantism provides a more charitable treatment than contextualism of metalinguistic judgements about the truth-value of knowledge attributions, comparative judgements about whether various knowledge attributions and denials conflict, and judgements about the correct use of ‘knows’ in speech and belief reports (e.g., Hawthorne 2004, §2.4–2.7; Williamson *forthcoming* and Wright *forthcoming*).¹ Here I focus just on the extent to which invariantism can offer a charitable treatment of intuitions about simple applications of ‘knows’ and related intuitions concerning assertion and practical reasoning. If an invariantist account need not treat these intuitions about contextualist cases as misguided, then the *prima facie* concern that invariantism is uncharitable does not arise.

In discussing this issue I will focus solely on so-called ‘classic invariantism’, which denies that the truth-value of attributions of knowledge depends on whether, in either the attributor’s or the subject’s context, the issue is important and error has been raised. Subject-sensitive invariantism also denies that the truth-value of knowledge attributions depends on the attributor’s context, but argues that it does depend on whether, in the subject’s context, the issue is important and error has been raised (e.g., Hawthorne 2004; Stanley *forthcoming*). In virtue of this subject-sensitivity, subject-sensitive invariantism can offer a charitable explanation of the contextualist’s intuitions at least in cases where the attributor is the subject. For instance, in a first-person contextualist case, the subject-sensitive invariantist can argue that, in the high context, it seems correct for the attributor to deny that she herself knows since the denial of knowledge would be false. Since the classic invariantist denies that the truth-value of knowledge attributions depends on whether, in either the subject’s or attributor’s context, the issue is important and error has been raised, it seems to face a tougher challenge in accommodating contextualist intuitions than subject-sensitive invariantism. In focusing on classic invariantism, then, we are making the task as

1. For counter-arguments see Cohen (*forthcoming*) and DeRose (*forthcoming*).

hard as possible for the invariantist. From now on, I will use the term ‘invariantism’ to refer to classic invariantism alone.

1. *Intuitions about contextualist cases*

In Cohen’s airport case, Mary and her companion John are wondering whether the New York flight stops in Chicago. Mary asks someone, Smith, whether he knows whether the flight stops there. Smith replies, ‘Yes, I know the New York flight stops in Chicago; my itinerary says so’. Suppose that Mary and John are in a low context (hereafter ‘LOW’), in which it’s not important whether the New York flight stops in Chicago, and no one has raised any error possibilities. In LOW, Mary would attribute knowledge that the New York flight stops in Chicago to Smith. Further, her attribution of knowledge to Smith seems correct. By contrast, it would seem incorrect if she were to deny that Smith knows. Further, it seems correct for her to assert that the New York flight stops in Chicago and to use that proposition in practical reasoning. For instance, she might reason that since the New York flight stops in Chicago, the journey will be longer than if the flight were non-stop and she should buy an extra magazine to read.

Now suppose that instead Mary and John are in a high context (hereafter ‘HIGH’) in which it’s very important that the plane stops in Chicago and John has raised the possibility of itinerary error (perhaps, Mary and John are heart surgeons due to perform an emergency operation in Chicago). In response to Mary’s query, Smith again responds ‘Yes, I know the plane is stopping in Chicago; my itinerary says so’. In this case, Mary would fail to attribute knowledge to Smith and her failure to do so seems correct. Instead, she would deny that Smith knows and this denial seems correct. Further, it would seem incorrect if Mary were to assert that the New York flight stops in Chicago, or to practically reason from this proposition. We may summarise the intuitions about the case as follows, letting ‘c’ stand for the proposition that the New York flight stops in Chicago:

Attribution intuition: in LOW, Mary would attribute knowledge that c to Smith and this seems correct; in HIGH, Mary would fail to attribute knowledge that c to Smith and this failure seems correct.

Denial intuition: in HIGH, Mary would deny that Smith knows that

c and this denial seems correct; in LOW, Mary would not deny that Smith knows that c and her failure to deny seems correct.

Assertion intuition: in LOW, it seems correct for Mary to assert that c; in HIGH, it seems incorrect for Mary to assert that c.

Practical reasoning intuition: in LOW, it seems correct for Mary to practically reason from c; in HIGH, it seems incorrect for Mary to practically reason from c.²

Contextualism provides a unified account of these data. The contextualist explains the attribution and denial intuitions by claiming that they reflect the truth-value of the relevant attributions and denials.³ On the contextualist view, in LOW, Mary's attribution, 'Smith knows that c', is true and so she would attribute knowledge to Smith and would not deny that he knows; further, her attribution and failure to deny seem correct. Further, the contextualist holds that, in HIGH, Mary's denial, 'Smith does not know that c', is true. As a result, she would fail to attribute knowledge to Smith and would instead deny that he knows; further, her failure to attribute and her denial seem correct. On the contextualist view, were Mary to self-attribute knowledge that c, the truth-value of her self-attributions would follow the same pattern as her attributions to Smith: in LOW it would be true for Mary to say 'I know that c' but, in HIGH, it would be false for her to say 'I know that c'. As we will see, the contextualist may use the truth-value of Mary's self-attributions to explain the assertion and practical reasoning intuitions by appeal to certain principles linking knowledge and appropriate assertion and practical reasoning.

According to the 'knowledge rule for assertion', or KR for short, one should assert that p only if one knows that p. Slote (1979), Unger (1975), Williamson (2000) and DeRose (2002) all endorse the knowledge rule.⁴

2. For discussion of the attribution and denial intuitions, see Cohen (1988) and DeRose (1992); for the assertion intuition, see DeRose (2002) and Hawthorne (2004); for the practical reasoning intuition see Hawthorne (2004) and Williamson (*forthcoming*). Although most contextualists and invariantists accept that the cases are characterised by these intuitions, some disagree. For instance, Pritchard (2005) rejects the denial intuition.

3. For a concern about the extent to which the truth-value of the relevant attributions/denials explains the intuitions, see Spicer (*forthcoming*).

4. DeRose (2002) also endorses the claim that that the knowledge rule is the only rule governing when a subject is well-enough positioned to assert that p, and uses this in an argument for contextualism. Brown (*forthcoming a*) argues against this claim and against the argument for contextualism based on it.

KR may be supported by noting that it explains the paradoxical nature of Moorean statements of the form ‘p but I don’t know that p’, and the fact that an assertion that p may be legitimately challenged by saying, ‘How do you know that p?’.

The contextualist may explain the assertion intuition by using KR in a contextualist form:

KR: one should assert that p only if ‘I know that p’ is true in one’s context.

According to contextualism, in HIGH it would be false for Mary to say ‘I know that c’. Thus, by KR, it is inappropriate for her to assert that c. By contrast, in LOW, it would be true for Mary to say ‘I know that c’, and so she meets the necessary condition for appropriately asserting p specified in KR. (Such an explanation of the assertion intuition is offered in DeRose 2002.)

The contextualist may use the truth-value of Mary’s self-attributions of knowledge to explain the practical reasoning intuition by exploiting a principle, KPR, which links practical reasoning and knowledge.

KPR: one should use p as a premise in practical reasoning if and only if ‘I know that p’ is true in one’s context.

KPR has been defended by Hawthorne and Williamson who argue that our intuitions about when a subject knows that p and when it’s appropriate for her to practically reason from p go together. For instance, before the lottery draw is announced it seems false for me to claim that I know my ticket is a loser. Further, it seems incorrect for me to rely on the proposition that my ticket is a loser in practical reasoning, say in the inference that since my ticket is a loser I should sell it for a penny. After the draw has been announced and my ticket is not the winning ticket, I do know that my ticket is a loser and it also seems appropriate for me to rely on that proposition in practical reasoning, say in the inference that since my ticket is a loser I should sell it for a penny (Hawthorne 2004, 30; see also Williamson *forthcoming*). The contextualist can exploit KPR to explain the practical reasoning intuition. In LOW, given that it would be true for Mary to say ‘I know that c’, by KPR it is appropriate for her to practically reason from c. In HIGH, since it would be false for Mary to say ‘I know that c’, by KPR it would

be inappropriate for her to practically reason from *c*. (Hawthorne 2004 argues that although contextualists can use KR and KPR to explain the assertion and practical reasoning intuitions, contextualists are committed to certain counterintuitive claims concerning practical reasoning and knowledge. I am setting aside consideration of such complex claims here and just examining how contextualists and invariantists deal with our simple intuitions about contextualist cases.⁵)

We have seen that, on the contextualist account, the attribution, denial, assertion and practical reasoning intuitions are correct. The attribution and denial intuitions reflect the truth-values of the relevant knowledge attributions; the assertion and practical reasoning intuitions reflect whether, given KR and KPR, it is appropriate for Mary to assert *c* or practically reason from *c*. Invariantism denies that the truth conditions of knowledge attributions depend on the attributor's context. So, it may seem that the invariantist must treat our intuitions about contextualist cases as misguided. For instance, suppose that the invariantist holds that, in HIGH, Smith knows that *c*. On that view, Mary's denial that Smith knows is literally false.

I will investigate whether invariantism must treat our intuitions about contextualist cases as misguided by examining a range of different non-sceptical invariantist views. In doing so, I will assume the principles KPR and KR. Some invariantists would prefer to offer a treatment of the contextualist intuitions while rejecting these principles (e.g., Pritchard and Rysiew reject KR⁶). However, these principles do receive intuitive support and are accepted by some invariantists, e.g., Williamson. Further, it is interesting to see if invariantists can offer a charitable treatment of the data while accepting the key assumptions of KR and KPR used in the contextualist treatment of the data.

5. Hawthorne argues that the contextualist is committed to the truth of such counterintuitive claims as 6) People often flat-out assert things that they do not know to be true but are not thereby subject to criticism; and 7) You should rely on propositions you don't know to be true in your practical reasoning (2004, 88).

6. Pritchard accepts the Gricean principle that one's assertions should be supported by adequate evidence (2005, 75, note 9); Rysiew endorses the claim that one should assert *p* only if one has a justified true belief that *p* (2001, 492).

2. *Pure error theory*

A pure error theory explains the contextualist data by appeal to mistaken judgements about knowledge, and without bringing in any further factors. A number of different mechanisms have been suggested by which HIGH may lead us and Mary to falsely suppose that she and Smith lack knowledge that *c*. Perhaps the salience of error in HIGH leads us and Mary to overestimate the likelihood of error via the availability heuristic (Hawthorne 2004), and/or it may be that we use convenient but highly fallible rules of thumb in determining whether someone knows, rules which lead us astray in HIGH (Williamson *forthcoming*). Suppose that whether Mary is in LOW or HIGH, it is true for Mary to say ‘Smith knows that *c*’ and ‘I know that *c*’, however Mary’s and our judgements do not reflect these epistemic facts. Focus first on Mary’s judgements about her own and Smith’s epistemic positions. Suppose that when Mary is in LOW, she judges truly that she and Smith know that *c*; however, when Mary is in HIGH, she judges falsely that she and Smith do not know that *c*. On this supposition, she would attribute knowledge to herself and Smith in LOW, but not in HIGH; instead, in HIGH, she would deny that she or Smith knows that *c*. Now focus on our judgements about Mary’s and Smith’s epistemic positions. Suppose that our judgements follow the same pattern as Mary’s; suppose that when we consider LOW, we judge truly that Mary and Smith know that *c*, but that when we consider HIGH, we falsely judge that neither Mary nor Smith know that *c*. As a result, in LOW it seems correct for Mary to attribute knowledge to herself and Smith, but in HIGH, it seems incorrect for Mary to do so; instead, in HIGH, but not LOW, it seems correct for Mary to deny that either she or Smith know. The pure error theorist could exploit our judgements about whether Mary knows combined with KR and KPR to explain the data concerning assertion and practical reasoning. By KR, one should assert that *p* only if one knows that *p*. If, in HIGH, we judge that Mary does not know that *c*, then we will also judge that it is incorrect for Mary to assert that *c*. If, in LOW, we judge that Mary knows that *c*, then we will also judge that she meets KR’s necessary condition for appropriately asserting that *c*. By KPR, one should use *p* as a premise in practical reasoning if and only if one knows that *p*. If, in HIGH, we judge that Mary does not know that *c*, then we will also judge that it is incorrect for Mary to practically reason from *c*. If, in LOW, we judge that Mary knows that

c, then we will also judge that it is appropriate for Mary to practically reason from c.

Notice that although the pure error theory explains the attribution, denial, assertion and practical reasoning intuitions, it treats them all as misguided. On the theory, these intuitions are based on Mary's and our false judgements about Smith's and Mary's epistemic state. For instance, when considering HIGH, we judge that it is incorrect for Mary to attribute knowledge to Smith since we falsely judge that Smith does not know that c. Similarly, considering HIGH, we judge that it is incorrect for Mary to assert c, or practically reason from c, for we mistakenly judge that Mary does not know that c. According to a pure error theory, when Mary is in HIGH, it is correct for her to say 'Smith knows that c' in the sense that it is literally true, and incorrect for her to say 'Smith does not know that c' in the sense that it is literally false. Assuming that the theory accepts KPR, it seems committed to saying that when Mary is in HIGH, since she knows that c, it is appropriate in the sense of KPR for her to practically reason from c. Further, if it accepts KR, then if, in HIGH, Mary knows that c, then she meets KR's necessary condition for appropriately asserting c. So while the pure error theory explains the attribution, denial, assertion and practical reasoning intuitions, it treats those intuitions as in fact incorrect.⁷ This motivates an examination of other invariantist accounts to see if they also treat our intuitions about contextualist cases as misguided.

3. *Bach and belief removal*

A different way to explain the contextualist intuitions is offered by the belief-removal model. At the heart of this view is the claim that HIGH leads the attributor to lose belief and so knowledge in the relevant proposition. On one possible development of the view, HIGH undermines belief by making error salient (for a discussion, see Hawthorne 2004, 169–170). On Bach's different development, HIGH provides the attributor with a practical reason to doubt the relevant proposition (Bach *forthcoming*). Here, I consider Bach's development of the view.

7. Note that a view which combines an error theory with further considerations need not treat all of our intuitions about contextualist cases as misguided. See the discussion of Bach's view and Williamson's view, §3 and §5.

Bach rejects the contextualist claim that the truth conditions of a knowledge attribution depend on the attributor's context. He defends a non-sceptical invariantist view on which, whether Mary is in LOW or HIGH, her attribution, 'Smith knows that c' is true. Bach argues that, in HIGH but not LOW, the high stakes give Mary reservations about the truth of c; she doubts that c. Given the high stakes, she has 'good practical reason' not to take Smith's word for whether the plane stops in Chicago.⁸ Mary's special practical interest 'raises the threshold for believing c', in other words she has practical reason to raise the level of epistemic justification required for belief; at these raised standards for belief, the itinerary evidence is not sufficient for belief. Bach argues that since Mary doubts that c it would be inappropriate for her to attribute knowledge of c to Smith. In general, an attribution of the form 'Smith knows that p' is true only if 'p' is true. So, if Mary were to assert 'Smith knows that c', Mary would commit herself to the truth of c. Since, in HIGH, she doubts that c, it would be inappropriate for her to attribute knowledge to Smith. In this way, Bach explains the attribution intuition.

By itself, the claim that, when Mary is in HIGH, Mary doubts that c does not explain why Mary goes so far as to deny that Smith knows that c. Mary's doubt about c may make inappropriate the attribution of knowledge to Smith but does not warrant the claim that Smith does not know that c. To explain the denial intuition, Bach further supposes that Mary's doubt that c leads her to conclude that her evidence doesn't epistemically justify her belief that c. His idea seems to be that Mary, finding herself doubting that c despite her evidence for c, would conclude that her evidence doesn't epistemically justify the belief that c. Since Smith has the same evidence as Mary, Mary would also judge that Smith's evidence does not justify his belief that c. Thus, Mary would go so far as to deny that Smith knows that c. For Bach, Mary's judgement that neither her nor Smith's evidence justifies the belief that c is mistaken.⁹

8. Discussing the high standards version of the bank case, Bach says that DeRose's wife 'has good practical reason not to take Keith's word for whether the bank is open on Saturday' (*forthcoming*, 25).

9. As Bach puts it, 'in the high-standards versions of the airport and bank cases, a special practical interest gives the attributor reservations about the truth of the proposition in question and raises the bar for attributing knowledge to someone else. This practical element is epistemically irrelevant, even if it can affect whether one can make the knowledge attribution' (*forthcoming*, 23).

In LOW and HIGH, Mary has the same evidence and has no positive reason to doubt that evidence. As an invariantist, Bach rejects the claim that the truth conditions of attributions of knowledge, or other epistemic states, such as justification, depend on the attributor's context. So, he holds that Mary has the same level of justification for *c* in LOW and HIGH. Since the evidence is sufficient for knowledge in LOW, it is also sufficient for knowledge in HIGH. Bach's idea, then, is not that in HIGH Mary has less justification for *c* than in LOW. Rather, he holds that in HIGH, but not LOW, Mary has a practical reason to doubt *c* which she confuses for the different claim that she lacks epistemic justification for believing that *c*.

There are, then, two elements to Bach's account. Like the pure error theorist, Bach holds that HIGH leads Mary to make a mistaken judgment about whether she and Smith know the relevant proposition. However, the account also contains an additional element, that HIGH leads Mary to lose belief in and so knowledge of the relevant proposition. This belief-removal element allows Bach, unlike the pure error theorist, to treat many of our intuitions about contextualist cases as correct. Although he holds that when Mary is in HIGH, it would be true for her to say 'Smith knows that *c*', it would nonetheless be inappropriate. For, in attributing knowledge to Smith Mary would be committing herself to something which, in HIGH, she doubts, namely *c*. Further, when Mary is in HIGH, Mary's failure to attribute knowledge to herself and her denial that she knows reflect the truth-values of the attribution and denial. If HIGH leads Mary to doubt that *c*, then she fails to meet the belief requirement for knowledge that *c*. So, it would be literally false for Mary to self-attribute knowledge and it is literally true for Mary to deny that she knows. So, Bach treats as correct the intuitions that, in HIGH, it is inappropriate for Mary to attribute knowledge to herself or Smith, and appropriate for her to deny that she knows.

Now consider the assertion and practical reasoning intuitions. If in HIGH, but not LOW, the stakes lead Mary to doubt that *c* and so fail to know that *c*, then Bach can exploit KPR and KR to argue that, in HIGH, it not only seems incorrect for Mary to practically reason from *c* and assert *c*, but it would in fact be incorrect for her to do so. By KPR and KR, it is appropriate for Mary to assert that *c* and practically reason from *c* only if she knows that *c*. Since, in HIGH, Mary fails to believe and so know that *c*, it would be inappropriate for her to assert that *c* or practically reason from *c*. Note that, in LOW, there is no similar

mechanism to dislodge Mary's belief that *c*, so, in LOW, Mary knows that *c*. As a result, it follows from KPR that, in LOW, it is appropriate for Mary to practically reason from *c*. Further, she meets KR's necessary condition for appropriately asserting *c*.

It seems, then, that the belief-removal element in Bach's account allows him, unlike the pure error theorist, to treat the attribution, assertion and practical reasoning intuitions as correct. Further, it allows him to treat as correct the first-person version of the denial intuition. Although Mary's lack of belief explains why her first-person denial of knowledge is correct, it does not help explain her third-person denial of knowledge to Smith. That Mary lacks belief, and so knowledge, that *c* does not explain her asserting that Smith does not know that *c*. Instead, Bach explains the denial intuition by appeal to the second element in his account, his idea that Mary's loss of belief leads her to falsely judge that she and Smith lack justification for and so knowledge that *c*. Note that Bach's explanation of Mary's denial that Smith knows treats this third-person denial as incorrect. On his view, in LOW and HIGH, Smith does have justified belief, and knowledge, that *c*. So, Mary's denial that Smith knows is false. It seems, then, that Bach's view provides a charitable treatment of all the intuitions about the airport case except the intuition that, in HIGH, it is correct for Mary to deny that Smith knows.

Since, with the exception of the third-person version of the denial intuition, Bach's account can treat all our intuitions about contextualist cases as correct and not misguided, it may seem a real rival to the contextualist's account. However, unfortunately, it faces several objections. The account is built around the idea that, when Mary is in HIGH, Mary's practical interest leads her to lose belief that, and so knowledge that, *c*. This directly explains why, in HIGH, it seems correct for Mary to fail to attribute knowledge that *c*, to assert *c* or use *c* in practical reasoning. Further, the loss of belief leads Mary to falsely believe that she and Smith do not have justification or knowledge that *c*. This false judgment explains why Mary denies that Smith knows, although this denial is literally false. However, it is not obvious that in all contextualist cases, the attributor loses the relevant belief in HIGH. For instance, consider a variant of the airport case in which in HIGH, it's pretty important that Mary and John get to Chicago for a business meeting, but not a matter of life or death. Further, so far no error possibilities have been explicitly raised in their conversation. They ask whether anyone knows whether

the New York flight is stopping at Chicago. Smith replies, ‘Yes, I know the flight stops there; my itinerary says so’. The following conversation then takes place:

John: Great. The plane’s going to Chicago, let’s take it.

Mary: Smith’s itinerary may say that it’s going to Chicago, but Smith doesn’t know that. The itinerary might be mistaken; perhaps there’s been a late change of schedule. We better check.

John: Well, it’s not very likely that the itinerary’s incorrect. We’ve certainly no reason to think there’s been a late change. How often have you come across an incorrect itinerary? Of course, the plane’s going to Chicago.

Mary: OK, I know it’s unlikely that the itinerary’s wrong. I believe the plane’s going to Chicago too. But that’s not the point. Just imagine what would happen if we took the flight and it doesn’t stop at Chicago. We’d miss the meeting. We can’t rule out the possibility of an itinerary error. So we don’t know the plane will stop and we better check.

Here both Mary and John make it clear that they believe that the New York flight stops in Chicago. Nevertheless Mary claims that they do not know that it does and her denial seems intuitively correct. But given that Mary and John do believe that the flight stops in Chicago, one cannot explain why it seems correct for Mary to deny that they know that the flight stops there by appeal to Bach’s account.

This example highlights a general worry facing belief-removal accounts.¹⁰ Whether HIGH leads one to doubt the relevant proposition is a matter of psychological fact. So, suppose instead that as a matter of psychological fact, HIGH does not lead Mary to doubt that *c*. It still seems inappropriate for Mary in HIGH to attribute knowledge that *c* to herself or Smith, to assert that *c* or to practically reason from *c*. But, if HIGH does not lead Mary to lose belief, it’s not clear how the

10. Compare Hawthorne’s objection to a range of contextualist and invariantist accounts which explain the contextualist data by appeal to the salience of error in HIGH. Hawthorne argues that such accounts, whether in the form of a classic invariantist belief-removal model, a subject-sensitive invariantist model or a contextualist model cannot explain why it is wrong, in HIGH, for Mary to assert that *p* and practically reason from *p* if Mary continues to dogmatically believe that *p* because, as a matter of psychological fact, the counter-possibilities are not salient to her (2004, 173–4).

belief-removal account can treat these intuitions charitably. The belief-removal model is non-sceptical and holds that, in LOW, Mary knows that *c*. Further, it is invariantist and holds that the attributor's context does not affect the truth conditions of attributions of knowledge and other epistemic states. Given this invariantism, if in LOW Mary knows that *c*, in HIGH, if Mary believes that *c*, she knows that *c*. If in HIGH Mary continues to believe and so know that *c* then the belief-removal model has no account of why, in HIGH, it is inappropriate for her to attribute knowledge to herself or Smith. Further, if in HIGH, Mary knows that *c* then, assuming KPR, it is appropriate for her to practically reason from *c*, and she meets KR's necessary condition for appropriately asserting that *c*. So if, in HIGH, Mary continues to believe that *c*, then the belief-removal model lacks an account of why it is wrong for Mary in HIGH to assert that *c*, practically reason from *c*, or attribute knowledge that *c*.

Note that the belief-removal account cannot overcome this difficulty even if it adds some mechanism independent of belief-removal by which HIGH leads Mary to falsely judge that she and Smith do not know that *c*. Suppose that HIGH does not lead Mary to lose belief that *c*, but on grounds independent of belief-removal, it leads her to judge falsely that she and Smith do not know that *c*. This false judgment would explain why Mary fails to attribute knowledge that *c*, denies knowledge that *c*, and fails to assert *c* and practically reason from *c*. However, so long as Mary still believes that *c* then on this invariantist view, she knows it. But, then, it would be literally true for her to self-attribute knowledge that *c*, appropriate in the sense of KPR for her to use it in practical reasoning, and she meets KR's condition for appropriately asserting that *c*. So, this modified view would still treat the attribution, assertion and practical reasoning intuitions as misguided.

Bach might try to overcome the problem of cases in which HIGH does not lead the attributor to lose belief in a different way. He could argue that even if, in HIGH, Mary continues to believe that *c*, her doxastic practices are criticisable. A praiseworthy agent would doubt that *c* given the stakes and so such an agent ought not to self-attribute knowledge that *c*, assert that *c*, or practically reason from *c*. Thus, he may say, even if Mary continues to believe that *c*, she is blameworthy for self-attributing knowledge that *c*, asserting *c* and practically reasoning from *c*. However, the fact that a different, and perhaps more praiseworthy agent who does not believe that *c* ought not to self-attribute knowledge,

assert *c* or practically reason from *c* does not show that Mary ought not to do so. For there is a relevant difference between them: Mary believes and so knows that *c*. As a result, her self-attribution is true; by KPR, it is appropriate for her to practically reason from *c*; and, further she meets KR's necessary condition for appropriately asserting that *c*.

A different reply might start from the thought that, given the stakes, Mary herself ought not to believe that *c* and so ought not to do things whose appropriate performance requires that she believe that *c*. However, on Bach's view, the relevant ought appears to be related solely to practical reasoning, and not epistemic justification. Since the belief-removal model is invariantist, it rejects the claim that whether a subject can be truly ascribed epistemic properties depends on how high the stakes are. So this view should hold that if, in HIGH, Mary continues to believe that *c*, then her belief has the same level of epistemic justification as it does in LOW. Given that, in LOW, Mary's belief that *c* counts as knowledge, so it does in HIGH. So, even if in HIGH Mary has a practical reason not to believe that *c*, she still has epistemic justification to believe that *c* and knows that *c*. So, in HIGH, if she continues to believe that *c*, she does know that *c*. As a result, were she to self-attribute knowledge, her attribution would be literally true. Further, since KR and KPR tie appropriate assertion and practical reasoning to Mary's epistemic state, not considerations of practical reasoning, if in HIGH Mary knows that *c* it is appropriate in the sense of KPR for her to practically reason from *c* and she meets KR's necessary condition for appropriately asserting that *c*.

It seems, then, that whether the belief-removal account can treat the intuitions about contextualist cases charitably depends on whether HIGH leads the attributor to lose belief. Whether HIGH does undermine belief is a matter of brute contingent psychological fact. We have seen that it is possible to construct contextualist cases in which HIGH does not undermine belief. For such cases, the belief-removal account cannot offer a charitable explanation of the data.

4. *Warranted assertibility manoeuvres*

A warranted assertibility manoeuvre, or WAM for short, makes no use of the idea that, when Mary is in HIGH, she doubts that *c* and so lacks knowledge that *c*. Rather, a WAM tries to explain the contextualist data

compatibly with accepting that, whether Mary is in HIGH or LOW, it would be true for her to say ‘Smith knows that c’ and ‘I know that c’. At the core of a WAM is the idea that the intuitions about contextualist cases can be explained by appeal to the truth-value of the propositions pragmatically conveyed by knowledge attributions, rather than the literal truth-value of those attributions. There are a number of different ways of filling out the basic strategy of a WAM. Here I will sketch Rysiew’s detailed example of such an account (see also Brown *forthcoming b*).

According to the contextualist, the strength of epistemic position required for knowledge depends on context. By contrast, on Rysiew’s invariantist position, there is a context-invariant level of epistemic position required for knowledge although, in some contexts, an attribution of knowledge may pragmatically convey that the subject has a stronger epistemic position. More specifically, he argues that a subject knows that p iff she can rule out the relevant alternatives to p, i.e., those alternatives which ‘we (normal) humans *take to be* the likely counterpossibilities to what the subject is said to know’ (Rysiew 2001, 488). However, even an irrelevant alternative may be salient in a conversation, say because it has been mentioned (Rysiew 2001, 488). In such a context, Rysiew claims that, via Grice’s rule of relevance, an attribution of knowledge pragmatically conveys that the subject can rule out the salient alternative(s). For instance, Rysiew argues that in both the low and high contexts of the airport case, it’s literally true for Mary to say ‘Smith knows that c’ or ‘I know that c’. For she and Smith can rule out the relevant alternatives to c. The possibility of an itinerary error is not a relevant alternative in either LOW or HIGH. However, in HIGH, the possibility of itinerary error is salient in the conversation so an attribution of knowledge would pragmatically imply that the subject of the attribution can rule out this possibility. Since neither Mary nor Smith can do so, the attribution is conversationally inappropriate, although literally true. As a result, Mary fails to attribute knowledge to either herself or to Smith. Instead, she makes the false assertion that she and Smith do not know, pragmatically conveying the truth that neither she nor Smith can rule out the possibility of itinerary error.

DeRose has raised doubts about the viability of such a WAM. Since I have discussed these arguments elsewhere, I merely summarise my conclusions here (see Brown *forthcoming a* and *forthcoming b*). DeRose (2002) presents an argument for contextualism and against an

invariantist WAM from ‘the knowledge account of assertion’, which combines KR with the claim that KR is the only rule governing how well-positioned a subject must be to warrantably assert *p*. Brown (*forthcoming a*) argues against DeRose’s claim that KR is the only such rule; here, we are merely granting KR. Earlier (1998) DeRose argued that a plausible WAM should exploit general conversational rules and should arise from a conflict of intuitions. As we have seen, Rysiew’s WAM does exploit a general conversational principle, namely Grice’s rule of relevance. Further, it arises from a conflict of intuitions. For instance, in the airport case, if we focus on the fact that Smith’s itinerary says that the flight stops in Chicago, it seems plausible that Smith knows, but if we focus on the importance of the issue to Mary and John and the salience of error, then it seems implausible to claim that Smith knows (Brown *forthcoming b*).

The most serious question DeRose raises about the viability of a WAM is his suggestion that a WAM cannot explain why, in HIGH, it seems correct to deny knowledge. According to the invariantist view suggested, the denial of knowledge is literally false but conveys a truth. However, DeRose argues that a literally false claim cannot seem correct even if it conveys a truth: ‘don’t we want to avoid falsehood both in what we implicate and (especially!) in what we actually say? So, it would seem that it would be unwarranted to assert a falsehood, even if doing so generates a true implicature’ (1998, 200). In response, Rysiew points out that speakers are often poor at distinguishing what’s literally said and what’s pragmatically conveyed. If speakers mistake what’s conveyed for what’s said then, contrary to DeRose, a falsehood which conveys a truth may seem correct. However, even if this is granted, it raises a further question. Under what circumstances are speakers likely to confuse the semantics and pragmatics of an utterance in such a way that they would take a literally false utterance which conveys a truth to be correct? Rysiew suggests that this is especially likely to happen when the utterance in question nearly universally conveys the relevant pragmatic claim (Rysiew 2000, 496). However, this explanation is problematic. The term ‘and’ nearly universally conveys ‘and then’ although that’s no part of its literal meaning. As a result, on Rysiew’s account, we should expect that speakers would confusedly read the conveyed temporal order into the literal meaning of a sentence containing ‘and’. If speakers did suffer this confusion, then in the situation in which Jack and Jill got married and *then* fell in love, we would expect speakers

to take the following utterance to be correct, even though it is literally false: ‘Jack and Jill did not fall in love and get married’. However, in the imagined circumstances, the utterance would not seem correct. Rather, it would seem appropriate to say ‘Jack and Jill fell in love and got married, but not in that order’.¹¹

It seems, then, that important work remains to be done to defend the explanation a WAM offers of the denial intuition. Setting that issue aside for now, let us examine the extent to which a WAM can treat as correct the contextualist’s intuitions. We have seen that a WAM explains how, in HIGH, the attribution of knowledge seems incorrect even though it is literally true by arguing that it conveys a falsehood. This explanation can treat the attribution intuition as correct, as long as we assume that the attribution intuition is not only concerned with whether the attribution is literally true. If the content of the intuitive judgement that the attribution is incorrect amounts to the claim that the attribution is literally false then, on a WAM, that intuition would be incorrect since, on this view, the attribution is literally true. However, it seems plausible that our judgements about whether a given assertion is correct concern not only whether it is literally true, but also whether it is conversationally relevant, e.g., whether the assertion would be misleading. So, it seems more appropriate to treat the content of the intuitive judgement that the attribution is incorrect as the claim that it is either false or conversationally inappropriate. On this understanding, the attribution intuition is correct, for the attribution of knowledge is true but would be misleading since it conveys a falsehood.

Notice that, if the content of our intuitive judgement that a knowledge attribution or denial is correct amounts to the claim that it is neither false nor conversationally inappropriate, then a WAM treats the denial intuition uncharitably.¹² A WAM explains the denial intuition by claiming that, in HIGH, the denial of knowledge seems correct although it is literally false because it conveys a truth. So, on this account, the judgement that, in HIGH, the denial is correct is misguided: even though the denial pragmatically conveys a truth, it is still inappropriate in the sense that it is literally false.

11. See Halliday (*forthcoming*).

12. The only way to avoid this consequence would be to argue, implausibly, that the content of our intuitive judgement that, in HIGH, it is correct to deny knowledge merely concerns whether the denial conveys a truth, and is not at all concerned with whether the denial is literally true.

Now consider how a WAM might be extended to explain the assertion and practical reasoning intuitions. It may be suggested that, in HIGH, a simple assertion seems incorrect for the same reason as an attribution of knowledge: although literally true, it conveys a falsehood. For instance, in the high context of the airport case, the possibility of itinerary error is salient. As a result, Rysiew argued that an attribution of knowledge that the flight stops in Chicago would pragmatically convey the falsehood that the subject can rule out the possibility of itinerary error. Similarly, it may be suggested that the mere assertion that the flight stops in Chicago would pragmatically convey that the subject can rule out the possibility of itinerary error. If a WAM can be extended in this way, then it would treat the assertion intuition as correct: it would be incorrect for Mary to assert that *c* since although the assertion is literally true in HIGH it would convey a falsehood.

The claim that, in HIGH, the simple assertion of *c* by Mary would convey that she can rule out the possibility of itinerary error seems plausible. However, someone who defends this claim about the pragmatic implications of simple assertions should relate this claim to a general theoretical account of pragmatic implications and how they are generated.¹³ However, Rysiew's main defence of his account of the pragmatic implications of knowledge attributions does not carry across to the case of simple assertion. Rysiew defends his claim that 'S knows that *p*' conveys that S can rule out the salient alternatives by endorsing Malcolm's view that "I know" is used in contrast with someone's (perhaps one's own) previous, present or potential, disbelief, or doubt or insecure belief' (Malcolm 1986, 212). Even if Malcolm's view applies to knowledge claims¹⁴, it is not plausible for simple assertions. Simple assertions are not typically used in contrast to disbelief or doubt. For instance, I may simply say to my colleague, 'It's 12 o'clock. I must go to my lecture', without any previous doubt or disbelief about what time it is.¹⁵

13. Pritchard defends this claim by appeal to an account of assertion which is a rival to KR, the claim that the assertion that *p* conveys that the speaker has adequate evidence for *p* (2005, 80). See also Brown (*forthcoming b*) for an account on which the assertion that *p* would be true but misleading in HIGH.

14. Halliday (*forthcoming*) casts doubt even on the view for the case of knowledge claims.

15. Rysiew further supports the account by appeal to an account of what's conveyed by an assertion which is a rival to the knowledge rule assumed here. On KR, the assertion that

If this first strategy of explaining the assertion intuition fails, then a WAM may employ a different strategy. We have seen that Rysiew's WAM suggests that Mary's denial of knowledge seems appropriate because, although it's literally false it conveys a truth, and we and Mary confuse the truth of what's conveyed for the truth of what's said. In other words, we and she take it that she does not know that *c*. By KR, it is appropriate for Mary to assert that *c* only if she knows that *c*. So, if we and she falsely judge that she does not know that *c*, we and she would also judge that it is inappropriate for her to assert that *p*.¹⁶ Similarly, by KPR, it is appropriate for Mary to use *c* in practical reasoning only if she knows that *c*. So, if we and she judge that she does not know that *c*, we and she would also judge that it is not appropriate for her to use *c* in practical reasoning. In this way, a WAM could be used to explain the intuitions concerning assertion and practical reasoning. Notice, though, that without further supplement, this explanation treats the assertion and practical reasoning intuitions as incorrect. On this view, in HIGH it is true though conversationally inappropriate for Mary to say 'Smith knows that *c*', or 'I know that *c*'. Given that and assuming KPR, it would be correct for Mary to practically reason from *c*, at least in the sense of KPR. Further, Mary meets the necessary condition for warranted assertion of *c* formulated in KR, namely that she knows that *c*. Without challenging KR and KPR, a WAM could avoid treating the assertion and practical reasoning intuitions as incorrect only by bringing in further elements, for instance further rules concerning practical reasoning. (In the next section, we see that Williamson brings in such extra rules to provide a charitable treatment of the assertion and practical reasoning intuitions.)

It seems, then, a WAM would treat many of the contextualist intu-

p conveys that the speaker knows that *p*. On Rysiew's rival view, a sincere assertion that *p* conveys that the speaker has a justified true belief that *p*. He argues that it follows from this account that 'I know that *p*' must convey more than that the speaker has a justified true belief that *p*. Instead, he suggests that it conveys that the speaker is confident or certain, that she can counter salient doubts about *p* (Rysiew 2001, 492). For otherwise, he claims, saying 'I know that *p*' would violate Grice's maxim of manner; it would be a less perspicuous way of conveying just what the simple '*p*' conveys. This second argument cannot be applied to argue that a simple assertion typically conveys that the speaker can counter salient doubts.

16. It is not obvious that Rysiew would accept this explanation of the assertion intuition since he holds that the sincere assertion that *p* conveys that the speaker has a justified true belief that *p*, which is a weaker condition than the condition that the speaker knows that *p* (Rysiew 2001, 492).

itions as misguided. Although a WAM can explain and treat as correct the attribution intuition, it would treat as incorrect the denial intuition and the practical reasoning intuition. Further, depending on which analysis it offers of the assertion intuition, it may treat that intuition as incorrect also. Given this result, it seems that we should search for a different invariantist response to contextualist cases, one which can treat more of our intuitions about such cases as correct.

5. *Williamson*

Williamson (*forthcoming*) rejects the idea that we should explain the contextualist data as a result of the Gricean principle of relevance. Instead, he focuses on the conditions for assertion and practical reasoning. He argues that, in some cases, appropriate assertion and practical reasoning require second-order knowledge. However, he claims that such second-order knowledge is lacking in contextualist cases. Williamson concentrates on first-person contextualist cases; however we will see that his account can be extended to third-person cases. To set up the account, let us start by focussing on Mary and her self-attribution of knowledge.

Williamson endorses a non-sceptical invariantist view on which, in both LOW and HIGH, Mary knows that *c* via Smith's testimony. This claim combined with the fact that Williamson also endorses both KR and KPR might seem to make it difficult for Williamson to explain the assertion and practical reasoning intuitions. By KPR, that a subject knows that *p* is the condition for her to appropriately use *p* in practical reasoning. By KR, a subject is warranted in asserting *p* only if she knows that *p*. So, on the view that, in both the LOW and HIGH contexts, Mary knows that *c*, in the sense of KPR and KR, it is appropriate for Mary to assert *c* and use *c* in practical reasoning. Williamson hopes to explain why it nonetheless seems incorrect for Mary to assert *c* or use *c* in practical reasoning by defending additional principles according to which in high stakes situations, assertion and practical reasoning require second-order knowledge and arguing that Mary lacks such knowledge (*forthcoming*, 230–33). In support of the latter claim, Williamson argues that although Mary knows that *c*, she only just counts as knowing; she is an example of a borderline case of knowledge. According to Williamson's anti-luminosity argument, for any borderline case of

a condition C, one does not know that C obtains, i.e. luminosity fails. So, although in both LOW and HIGH, Mary knows that c, she does not know that she knows that c (*forthcoming*, 232). Further, Williamson claims that, when the stakes are high, a subject is blameworthy for asserting p or practically reasoning from p if she does not know that she knows that p. His general idea is that when the stakes are high, the question of whether the subject has warrant for asserting p and practically reasoning from p becomes pressing. In high stakes cases, we are not merely interested in whether the subject meets the condition for appropriate assertion and practical reasoning, but also in whether she knows that she does, in other words whether she knows that she knows that p. For instance, he argues that when the stakes are high, a subject is blameworthy for using p in practical reasoning if she knows p but does not know that she knows that p. Similarly, when the stakes are high, a subject is blameworthy for asserting p if she knows p but does not know that she does. As a result, in HIGH but not LOW, it seems incorrect for Mary to assert c or practically reason from c.

Mary's lack of second-order knowledge that c also explains why, in HIGH, it seems correct for her to fail to attribute knowledge to herself. By KR, it is appropriate for Mary to assert 'I know that c' only if she knows that she knows that c. But, she lacks such second-order knowledge. The explanation of why, in HIGH, it seems correct for Mary to fail to attribute knowledge to herself can be extended to explain why, in HIGH, it seems correct for Mary to fail to attribute knowledge to Smith. Like Mary, Smith counts as a borderline case of knowing that c. The conclusion of Williamson's anti-luminosity argument is that, for any borderline case of a condition C, one does not know that C obtains. Thus, Mary does not know that Smith knows that c. By KR, it would be appropriate for Mary to say 'Smith knows that c', only if Mary knows that Smith knows that c. But Mary lacks that second-order knowledge about Smith.

Although Mary's lack of second-order knowledge explains why, in HIGH, it seems correct for Mary to fail to attribute knowledge that c, assert c and practically reason from c, it does not explain why it seems correct for her to deny that she and Smith know that c. The inappropriateness of attributing knowledge to herself or to Smith does not show that it is appropriate to deny that either she or Smith know that c. Instead, Williamson explains the denial intuition by appeal to psychological bias. He argues that, when Mary is in HIGH, the salience of error leads Mary

to judge falsely that she and Smith do not know that *c*. As a result, Mary denies that she and Smith know that *c*. Further, when we consider Mary in HIGH, the salience of error leads us to judge falsely that neither she nor Smith knows that *c*. As a result, it seems correct for Mary to deny that she and Smith know that *c* (*forthcoming*, 234–5).

As we have seen, a failure of second-order knowledge is at the core of Williamson's account of the assertion, practical reasoning and attribution intuitions. There may be some contextualist cases in which the subject not only has first-order knowledge but also has second-order knowledge. Williamson suggests that for such cases, he may apply his explanation of the assertion, practical reasoning and attribution intuitions at a higher level of knowledge. For finite minds like ours, knowledge does not iterate indefinitely. A subject who knows that *p* and knows that she knows that *p*, will fail to know at some higher level. When the stakes are high enough, one may be blameworthy for asserting *p*, practically reasoning from *p*, or self-attributing knowledge that *p* if, although one knows that one knows that *p*, one lacks some higher level of knowledge (*forthcoming*, 234).¹⁷

Having seen how Williamson explains the contextualist data, let us assess how charitable his treatment is. For reasons which are by now familiar, Williamson's treatment of the denial intuition is uncharitable. On his view, Mary goes so far as to deny that Smith knows since she falsely believes that Smith lacks knowledge. It is less clear whether Williamson treats the other intuitions charitably since his account effectively introduces several different notions of correctness for assertion and practical reasoning. The extent to which Williamson's account is charitable turns out to depend on which of these notions of correctness we take our intuitions to be concerned with. Williamson endorses KR and KPR according to which knowledge that *p* is required for appropriate assertion and practical reasoning. Let us use the term 'appropriate*' for the notion of appropriateness at play in these two rules. As we have seen, Williamson introduces a further dimension of assessment of assertion and practical reasoning, whether the assertion or reasoning is blame free. He holds that, in high stakes cases, a subject is blameworthy for asserting that *p* or practically reasoning from *p* unless she knows that she knows that *p*. An assertion or piece of reasoning might be appropri-

17. For some concerns about Williamson's approach to contextualist cases in which the attributor has second-order knowledge, see Brown (*forthcoming c*).

ate* while also blameworthy, if the subject knows that *p*, but does not know that she does.

It turns out that the distinction between the notions of being appropriate* and being blame free makes no difference to whether Williamson offers a charitable treatment of the intuition that, in HIGH, it is incorrect for Mary to attribute knowledge. Were Mary to attribute knowledge, her attribution would be inappropriate* and blameworthy. According to KR, it is appropriate* for Mary to assert that she (or Smith) knows that *c* only if she knows that she (or Smith) knows that *c*. Since she lacks such second-order knowledge, the attribution would be inappropriate*. According to Williamson's further principle, if the stakes are high, a subject is blameworthy for asserting *p* if she doesn't know that she knows that *p*. So, given the stakes, Mary's assertion that she (or Smith) knows that *c* would be blameworthy unless she knows that she knows that she (or Smith) knows that *c*. Since Mary lacks second-order knowledge that she (or Smith) knows that *c*, she lacks the third-order knowledge required for her attribution of knowledge to be blame free.

By contrast, the distinction between being appropriate* and blame free is important for determining whether Williamson offers a charitable treatment of the assertion and practical reasoning intuitions. When Mary is in HIGH, it seems incorrect for her to assert that *c* or practically reason from *c*. Williamson attempts to explain this intuition by arguing that although, in HIGH, it would be appropriate* for Mary to assert that *c* or practically reason from *c* (since she knows that *c*), she would be blameworthy for doing so (since she does not know that she knows that *c*). Whether this explanation offers a charitable explanation of the intuitions depends on the content of those intuitions and, in particular, in what sense(s) we judge that it would be incorrect for Mary in HIGH to assert that *c* or practically reason from *c*. If the content of our judgement is that it is inappropriate* for Mary to assert that *c* or practically reason from *c*, then the judgement is false. For, on Williamson's account, it is appropriate* for Mary to do these things. Alternatively, the content of our judgement may be that it is blameworthy for Mary to assert that *c* or practically reason from *c*. On this suggestion, our intuition is correct for, on Williamson's account, Mary would be blameworthy for asserting that *c*, or practically reasoning from *c*. Instead of focussing on just one of these notions of correctness, our intuitive judgement may be concerned with both the notions of being appropriate* and blame free. On this view, the judgement that it would be incorrect for Mary

to assert that *c*, or practically reason from *c*, should be interpreted as the claim that it is either inappropriate* or blameworthy for her to do so. Notice that this last view need not be committed to the claim that subjects explicitly distinguish the two notions, but rather that their judgements of correctness factor in both notions, even if they do not clearly distinguish them. On this third interpretation, Williamson can treat the practical reasoning and assertion intuitions as correct: although it would be appropriate* for Mary to assert that *c*, or practically reason from *c*, it would be blameworthy for her to do so. A last suggestion might be that our intuitions of incorrectness are vague and do not clearly focus on any one notion of correctness, or any specific combination of notions. On this view, our intuition that, in HIGH, it is incorrect for Mary to assert or practically reason from *c*, would turn out to lack a determinate truth-value.

We have seen that Williamson can offer a charitable treatment of the intuition that, in HIGH, it is incorrect for Mary to attribute knowledge that *c*, assert that *c*, or practically reason from *c*. He can do so by defending either of two claims, namely that the content of the intuitive judgement of incorrectness is concerned only with the notion of blameworthiness, or that it factors in both the notions of blameworthiness and appropriateness as suggested in the third interpretation above. Of course, in assessing how charitable Williamson's account is it is not only important to see whether he can offer a charitable treatment of our intuitions concerning HIGH, but also whether he can offer such a treatment of our intuitions concerning LOW. It seems easy for him to offer a charitable treatment of the intuition that, in LOW, it is correct for Mary to assert that *c* or practically reason from *c*. In LOW, Mary's assertion and practical reasoning are not only appropriate* (since she knows that *c*), but also blame free. For, it is only in high stakes cases that one is blameworthy for asserting that *c* or practically reasoning from *c* without knowing that one knows that *c*. So, under any of the first three interpretations offered above of our intuition that it is correct for Mary in LOW to assert that *c*, or practically reason from *c*, the intuition turns out to be correct.¹⁸

18. (It may be argued that the relevant intuition is correct even on the last interpretation on which our intuitive judgements are vague. Whether those judgements are precisified to the notion of being appropriate*, being blame free, or, neither inappropriate nor blame free, they count as correct. If the judgement is correct on any precisification, then perhaps it can be treated as correct *tout court*.)

The varying potential interpretations of the content of our intuitions about Mary do affect whether Williamson can offer a charitable treatment of the intuition that, in LOW, it seems correct for Mary to attribute knowledge. On Williamson's view, in LOW and HIGH, Mary and Smith know that *c*, but Mary lacks second-order knowledge of this fact, she fails to know that she and Smith know that *c*. As a result, by KR, in LOW, it is inappropriate* for Mary to attribute knowledge to herself or to Smith, even though the attributions are literally true. This might initially seem to make it difficult for Williamson to offer a charitable treatment of the intuition that, in LOW, it is correct for Mary to attribute knowledge.

Williamson may reply to this worry by pointing out that, on his view, it is not always a great crime if KR is broken. He appeals to this idea in defending KR against the objection that we do not always take great pains to verify a proposition before asserting it, e.g., in gossip or casual conversation (2000, 258–59). Williamson argues that this does not show that some assertions are not governed by KR, but rather that it is not always a serious matter if one breaks this rule. As he puts it, 'when we are relaxed in applying the rule, we feel entitled to assert *p* whenever we are not confident that we do not know *p*' (2000, 259). Similarly, in LOW where not much is at stake, we feel that it is acceptable for Mary to attribute knowledge even though she is breaking KR in doing so. In this way, Williamson can explain why it seems intuitively correct to attribute knowledge in LOW.

Even if Williamson can explain why it seems correct for Mary to attribute knowledge in LOW, it is a further question whether he can treat this intuition as correct. We earlier distinguished three different ways of interpreting the content of our intuitive judgements about contextualist cases. On the first interpretation, the content of our judgement of intuitive correctness is that Mary's attribution is appropriate*. On this interpretation, our intuitive judgement is incorrect for Mary's attribution of knowledge breaks KR. On the second interpretation, the content of our judgement of intuitive correctness is that Mary's attribution is blame free. On this second interpretation, our judgement is correct for, given the low stakes, she's not blameworthy for breaking KR. On the third interpretation, the content of the relevant judgement is that Mary's attribution is neither inappropriate* nor blameworthy. Notice that, on this third reading, the intuition is incorrect, for although Mary isn't blameworthy for breaking KR, she is nonetheless breaking it. It turns

out, then, that there is an asymmetry between the intuitions concerning LOW and HIGH. Earlier, we saw that Williamson could provide a charitable treatment of the HIGH intuitions on either the interpretation that the content of our intuitive judgements of correctness solely concerns the notion of being blameworthy, or the interpretation that it concerns the notion of being neither blameworthy nor inappropriate. However, we have now seen that a charitable treatment of all the intuitions concerning LOW requires the interpretation on which the content of our intuitive judgements of correctness solely concerns the notion of being blameworthy.¹⁹

We have seen, then, that on a suitable interpretation of our intuitive judgements concerning contextualist cases, Williamson can offer a charitable treatment of all these judgements with the exception of the denial intuition, namely the intuition that, in HIGH, it is correct for the attributor to deny knowledge. Since Williamson holds that the judgement that, in HIGH, it is correct to deny knowledge is the result of a false judgement, he treats the denial intuition as misguided. It seems interesting that Williamson is not alone in treating the denial intuition as misguided. Rather, it is a common feature of all the invariantist accounts examined that they provide an uncharitable treatment of the denial intuition. This is not accidental. The accounts examined are all non-sceptical invariantist accounts. On a non-sceptical invariantist view, in LOW, the subject knows the relevant proposition. Given invariantism if, in HIGH, the subject still believes that proposition, her belief constitutes knowledge, and so the denial of knowledge is literally false. As a result, so long as it is part of the content of the HIGH judgement that the denial is correct, that the denial is not false, then the denial intuition is misguided. On this understanding of the denial intuition, the only way for a non-sceptical invariantist account to avoid treating the denial intuition as incorrect would be by arguing that HIGH leads the subject of that denial to lose the relevant belief, and so knowledge.

19. Notice that Williamson's account can explain what is wrong with a dogmatist who, in HIGH, asserts that *c*, practically reasons from *c*, or attributes knowledge that *c*. Suppose that despite the high stakes in HIGH, Mary continues to believe that she and Smith know that *c*, and that she asserts that *c*, practically reasons from *c*, and attributes knowledge that *c* to herself and Smith. Given Williamson's account of practical reasoning and assertion, since Mary does not know that she knows that *c*, it follows that, in HIGH, she would be blameworthy if she asserted that *c* or practically reasoned from *c*, and it would be inappropriate* and blameworthy if she attributed knowledge that *c* to herself and Smith.

Even, then, that would only enable us to treat as correct first-person denials of knowledge where the subject is the attributor, but not third-person denials in which the subject is distinct from the attributor. For even if the attributor's being in HIGH undermines her, the attributor's, belief, it needn't undermine the belief of distinct subjects to whom she may attribute knowledge, subjects who themselves need not be in a high stakes context. It seems, then, that no non-sceptical invariantist account can treat as correct the denial intuition in both first and third-person versions.

6. *Conclusion*

Contextualists support their account of the intuitions concerning contextualist cases by claiming that it provides a more charitable treatment of these intuitions than the invariantist account. Initially, it may seem difficult to see how the invariantist could offer a charitable treatment of the data since she denies that the truth conditions of knowledge attributions depend on the attributor's context. However, of the variety of invariantist accounts considered here, it turns out that only the pure error theory offers an uncharitable treatment of all the contextualist intuitions. A WAM can charitably treat at least the attribution and perhaps the assertion intuitions, although it provides an uncharitable treatment of the denial and practical reasoning intuitions. Further, both the belief-removal account and Williamson's luminosity account can provide a charitable treatment of all the contextualist intuitions with the exception of the denial intuition. So it is far from obvious that charity considerations clearly support contextualism over invariantism.²⁰

20. Thanks to helpful comments from participants at the Amsterdam conference on Epistemological Contextualism.

REFERENCES

- Bach, K. (1994). 'Conversational Implicature', *Mind and Language* 9, 124–62.
- (forthcoming). 'The Emperor's New 'Knows'', *Contextualism in Philosophy: On Epistemology, Language and Truth*, (eds.) Gerhard Preyer and Georg Peter, Oxford University Press, Oxford. (Page references from the manuscript.)
- Brown, J. (forthcoming a). 'Adapt or Die: the Death of Invariantism?', *Philosophical Quarterly*.
- (forthcoming b). 'Contextualism and Warranted Assertibility Manoeuvres', *Philosophical Studies*.
- (forthcoming c). 'Williamson on Luminosity and Contextualism', *Philosophical Quarterly*.
- Cohen, S. (1988). 'How to be a Fallibilist', *Philosophical Perspectives* 2, 91–123.
- (1999). 'Contextualism, Scepticism, and the Structure of Reasons', *Philosophical Perspectives* 13, 57–89.
- (2000). 'Contextualism and Scepticism', *Philosophical Issues* 10, 94–107.
- (forthcoming). 'Knowledge, Speaker and Subject', *Philosophical Quarterly*.
- DeRose, K. (1992). 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research* 52, 913–29.
- (1995). 'Solving the Sceptical Problem', *Philosophical Review* 104, 1–52. Page references from *Scepticism*, (eds.) K. DeRose and T. Warfield (1999), 183–219, Oxford University Press, Oxford.
- (1998). 'Contextualism: an Explanation and Defense', *The Blackwell Guide to Epistemology*, (eds.) J. Greco and E. Sosa, 187–206, Blackwell, Oxford.
- (2002). 'Assertion, Knowledge and Context', *Philosophical Review* 111, 167–203.
- (forthcoming). 'The Ordinary Language Basis for Contextualism', *Philosophical Quarterly*.
- Halliday, D. (forthcoming). 'What Explains Our Intuitions About Knowledge Ascriptions?', *Proceedings of the Aristotelian Society*.
- Hawthorne, J. (2004). *Knowledge and Lotteries*, Clarendon Press, Oxford.
- Malcolm, N. (1986). *Wittgenstein: Nothing is Hidden*, Blackwell Publishers, Cambridge, MA.

- Pritchard, D. (2005). *Epistemic Luck*, Oxford University Press, Oxford. (Page references from the pre-publication manuscript.)
- Rysiew, P. (2001). 'The Context-Sensitivity of Knowledge Attributions', *Nous* 35, 477–514.
- Schaffer, J. (2005). 'Scepticism, Contextualism and Discrimination', *Philosophy and Phenomenological Research* 69, 138–155.
- Slote, M. (1979). 'Assertion and Belief', *Papers on Language and Logic*, (ed.) J. Dancy, Keele University Library, Keele.
- Stanley, J. (*forthcoming*). 'Context, Interest-Relativity and Knowledge', *Philosophical Studies*.
- Spicer, F. (*forthcoming*). 'Epistemic Intuitions and Epistemic Contextualism', *Philosophy and Phenomenological Research*.
- Unger, P. (1975). *Ignorance*, Oxford University Press, Oxford.
- Williamson, T. (2000). *Knowledge and Its Limits*, Oxford University Press, Oxford.
- (*forthcoming*). 'Contextualism, Subject-Sensitivity Invariantism, and Knowledge of Knowledge', *Philosophical Quarterly*.
- Wright, C. (*forthcoming*). 'On the Contextualist Way with Scepticism', *Philosophical Quarterly*.