

Refined Invariantism

Jacques-Henri Vollet (University of Geneva)

To appear in *Theoria: A Swedish Journal of Philosophy*

Abstract A certain number of cases suggest that our willingness to ascribe ‘knowledge’ can be influenced by practical factors. For revisionary proposals, they indicate that the truth-values of ‘knowledge’ ascriptions vary with practical factors. For conservative proposals, on the contrary, nothing surprising is happening. Standard pragmatic approaches appeal to pragmatic implicatures and psychological approaches to the idea that belief formation is influenced by practical factors. Conservative proposals have not yet offered a fully satisfactory explanation, though. In this paper, I introduce and defend a third conservative proposal which I call ‘Refined Invariantism’. The two main claims of this proposal are that (1) we should distinguish between high stakes cases in which the subject does not believe (that he knows) the target proposition and those in which he believes (that he knows) the target proposition and that (2) we should adopt a psychological treatment for the first kind of case and a pragmatic treatment based on the epistemic standards for appropriate assertion and action for the second kind of case. I argue that this new combined approach avoids the main pitfalls of its two conservative rivals and that it gives new life to the generality objection leveled against revisionary views.

1. Introduction

A certain number of cases discussed in the philosophical literature show that our willingness to ascribe or deny ‘knowledge’ can be influenced by practical factors (such as the cost of being wrong). Here are what are probably the best-known cases, proposed by Keith DeRose (1992: 913):

Bank Case A (LOW). My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on

Saturday morning. My wife says, “Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays.” I reply, “No, I know it'll be open. I was just there two weeks ago on Saturday. It's open until noon.”

Bank Case B (HIGH). My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and very important check. If our paychecks are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a *very* bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, “Banks do change their hours. Do you know the bank will be open tomorrow?” Remaining as confident as I was before that the bank will be open then, still, I reply, “Well, no. I'd better go in and make sure.”

While it seems appropriate for Keith to assert ‘I know that the bank will be open tomorrow’ in LOW, it also seems appropriate for him to assert ‘I don't know that the bank will be open tomorrow’ in HIGH. Yet, LOW and HIGH differ only in the two following ways. First, it is very important for Keith and his wife that they deposit their paychecks on or before Saturday in HIGH, but that's not important in LOW. Second, Keith's wife makes salient a possibility of error in HIGH which is not salient in LOW.

Philosophers have offered very different diagnoses about these cases. According to revisionary proposals we should grant the surprising claim that the truth-values of ‘knowledge’ ascriptions and denials are determined by practical factors. The two main revisionary proposals are Epistemic Contextualism – the view according to which the meaning of ‘know’ is sensitive to what is practically at stake (among other things) – and Subject-Sensitive Invariantism – the view according to which there is a practical condition on knowledge, e.g. that S knows that p only if it is rational for S to act as if p.¹ If this kind of view is correct, we can say that Keith says something true when he says ‘I know that the bank

¹ Contextualism is defended in particular by Cohen (1999), DeRose (1992; 2009), Blome-Tilmann (2014), Ichikawa (2017), and Subject-Sensitive Invariantism by Hawthorne (2004), Stanley (2005), Fantl and McGrath (2009). A third revisionary proposal is relativism, which is defended by McFarlane (2011).

will be open' in LOW, and that he says something true when he says 'I do not know that the bank will be open' in HIGH.

On the other hand, standard conservative (non-sceptical) approaches have proposed two other diagnoses.² According to standard pragmatic approaches, Keith says something true when he says 'I know that the bank will be open' in LOW, and he says something false when he says 'I do not know that the bank will be open' in HIGH, but asserting this literal falsehood is appropriate due to a pragmatic implicature.³

According to psychological approaches, Keith says something true when he says 'I know that the bank will be open' in LOW and he says something true when he says 'I do not know that the bank will be open' in HIGH. But this is explained by the very plausible assumption that practical factors influence knowledge and belief indirectly by influencing the belief forming mechanisms. In particular, when high stakes or possibilities of error are perceived, it becomes more difficult to form a settled belief. We can then say that Keith does not know in HIGH because he does not believe the target proposition (in the sense relevant for knowledge), and that he does not believe this proposition because his belief forming mechanisms are somewhat affected by the consideration of the possibility of error and the stakes.⁴

It is worth emphasizing how the two conservative approaches differ. In focusing exclusively on what is appropriate for Keith to say in HIGH, standard pragmatic approaches implicitly suggest that being in a high stakes situation is not particularly relevant to whether Keith can know the target proposition. In focusing exclusively on whether Keith believes the target proposition in HIGH, psychological approaches share with revisionary approaches the supposition that being in a high stakes situation is crucial to whether one can know the target proposition (even if these approaches differ on how and why that is the case).

In this paper, I focus on the main conservative proposals. I explain why they seem ultimately unsatisfactory. However, the consideration of their limits will suggest a new view which I call 'Refined Invariantism'. This view takes seriously the distinction between high stakes cases in which the subject does not believe (that he knows) the target proposition and high stakes cases in which the subject believes (that he knows) the target proposition. Its core insight is that two

² In this discussion, a view is classified as 'non-sceptical' if it maintains that Keith knows in LOW. For sceptical approaches to the problem, see Davis (2007) and Fassio (forth.). I will leave these approaches aside.

³ Advocates of standard pragmatic approaches include Rysiew (2001; 2005; 2007), Brown (2006), Hazlett (2009), Pritchard (2010), Lutz (2014) and Gerken (2017).

⁴ For psychological approaches, see Bach (2005; 2008; 2010), Williamson (2005), Nagel (2008, 2010a, 2010b), Engel (2009).

different explanations should be given for these two different cases: a psychological explanation for the first kind of case and a *non-standard* pragmatic explanation appealing to the epistemic standards of appropriate assertion and action for the second kind of case. As we will see, this new conservative approach gives new life to the generality objection which has been raised against revisionary views, in particular against epistemic contextualism. It also avoids the pitfalls of rival conservative approaches.

My plan is as follows. In section 2, I consider the standard pragmatic approaches. In section 3, I consider the two main psychological approaches. In section 4, I propose and defend my new conservative approach.

2. Standard pragmatic approaches

In general, pragmatic approaches seek to explain the variability exhibited in cases like LOW and HIGH by focusing on potential pragmatic effects of ‘knowledge’ attributions (and denials). In this section, I consider the *standard* pragmatic approaches. I briefly review the five main pressing problems that they face.

According to standard pragmatic approaches, Keith says something true when he says ‘I know that the bank will be open’ in LOW and he says something false when he says ‘I do not know that the bank will be open’ in HIGH. However, we are told, asserting this literal falsehood in HIGH is appropriate for this generates an implicature which is appropriate and true. For some writers, what is implicated is that Keith is not in a position to eliminate the salient alternative that the bank has changed its hours.⁵ For other theorists, what is implicated is that Keith is not in a position to act on the proposition that the bank is open.⁶ Yet other philosophers think that what is implicated is a recommendation not to come back but to rather deposit the pay check right now.⁷ In addition, standard pragmatic approaches maintain that it would be inappropriate although true for Keith in HIGH to say ‘I know that the bank is open’, for this would pragmatically communicate something inappropriate and false (for example, that Keith is in a good enough epistemic position to act on the target proposition).

These pragmatic approaches rely on the distinction between the literal meaning of an utterance (“what is said”) and its pragmatic meaning (“what is conveyed”). To illustrate this distinction, suppose I ask you whether John is a good philosopher and you tell me that he has

⁵ See Rysiew (2001).

⁶ See Brown (2006).

⁷ See Gerken (2017).

good handwriting. I will presumably conclude that you don't think John is a good philosopher. Although that's not what you literally say, that's what you pragmatically communicate. Following Grice (1989), we can understand the pragmatic meaning of an utterance as what is conveyed by this utterance in virtue of the literal meaning of the uttered sentence, the context (including the common background knowledge) and general conversational rules.

The distinction between pragmatic and semantic meaning seems clear, but many critics have put forth reasons to doubt that it can be put to work to provide an adequate explanation of the bank cases. Let us begin by clarifying what the explanatory challenge about the bank cases is supposed to be.

As said above, it seems that:

- (1) Keith's 'knowledge' denial is appropriate in HIGH.

It also seems that:

- (2) Keith's 'knowledge' ascription is inappropriate in HIGH.

Further, many philosophers (but not all) accept that:

- (3) Keith's 'knowledge' denial *seems* true in HIGH.
- (4) Keith's 'knowledge' ascription *seems* false in HIGH.

Finally, virtually everyone accepts that:

- (5) Keith can rationally act on the proposition that the bank is open in LOW.
- (6) Keith cannot rationally act on the proposition that the bank is open in HIGH.

Advocates of standard pragmatic approaches can easily explain claims (5) and (6) if they deny that knowledge is the norm of action and if they opt for a variable epistemic norm instead.⁸ On this view, in LOW, Keith satisfies the epistemic norm for appropriate action on the proposition that the bank is open, whereas, in HIGH, Keith does not satisfy this epistemic norm. Claims (5) and (6) do not seem to pose a special problem for standard pragmatic approaches.

Claims (3) and (4) are more controversial than claims (1) and (2). They are accepted by prominent proponents of standard pragmatic approaches.⁹ It seems clear, at least, that the attraction of these pragmatic approaches will be considerably weakened if they cannot explain (3) and (4), for it seems somewhat difficult to deny these intuitions.¹⁰

⁸ For such proposals, see among others Brown (2008); Levin (2008); Locke (2015); Gerken (2017).

⁹ For example, Brown (2006: 410).

Although initially promising, standard pragmatic approaches encounter at least five pressing problems. The first problem concerns (3) and (4). Advocates of standard pragmatic approaches must argue that literally false assertions (e.g. “I do not know that the bank is open”) may seem true in virtue of true implicatures (e.g. “I’m not in a good enough epistemic position to act”) and that literally true assertions (e.g. “I know that the bank is open”) can seem false in virtue of false implicatures (e.g. “I’m in a good enough epistemic position to act”).¹¹ Although advocates of standard pragmatic approaches often claim that we sometimes confuse the pragmatic and the semantic meaning, or that we often directly and exclusively focus on the pragmatic meaning,¹² that does not seem to be what happens for most cases of implicatures.¹³ There might be some cases in which false implicatures make literally true utterances appear false, but proponents of standard pragmatic approaches must still provide an explanation of why that should be the case when it comes to implicatures generated by ‘knowledge’ ascriptions (and denials).¹⁴ Call this problem the ‘apparent truth-value problem’. A second pressing problem concerns (1). According to standard pragmatic proposals, it’s appropriate for Keith to say something false (‘I do not know’) in HIGH in order to convey something true (e.g. that he is not in a good enough epistemic position to act). But, in general, it is doubtful that it can be appropriate to say (and convey) something false in order to convey

10 See Dinges (2018: sect. 3). Dimmock and Huvenes (2014: fn 6) grant that we might want to reject (4) and Hazlett (2009: 612), Pritchard (2010: 88) and Lutz (2014) reject (3). As we will see in section 4, a crucial claim of the proposal defended in this paper is that we should draw a distinction between two kinds of high stakes cases: one in which the subject confidently believes (that he knows) the target proposition, and one in which the subject does not confidently believe (that he knows) the target proposition. This approach can explain why we may have divergent intuitions about what seems true or false in HIGH, which I take to be a point in favour of the proposed approach.

11 See Dinges (2018: sect. 1).

12 See, e.g., Rysiew (2001: 486-487; 2007: 640), Brown (2006: 428) and Hazlett (2009: 610).

13 Dinges (2018: sect. 3) shows that this is not what happens at least in most cases of “additive” implicatures, in which what is conveyed is the literal meaning plus the implicatures (like in Grice’s famous petrol case). In particular, if the apparent truth-value is the result of the truth-value of the literal meaning in addition to the truth-value of the semantic meaning, (additive) true implicatures cannot make a literally false claim appear true, since the global seeming truth-value will be the falsity. Further, Dinges (2018) argues that if there is an implicature in the bank cases, it is additive and not “substitutional” (what is conveyed is not merely the pragmatic meaning but also the semantic meaning of the utterance). Substitutional implicatures require that, given the common ground, what is said is incompatible with the speaker being seen as observing (or trying to observe) the general conversational rules. In the bank cases, given the common ground (and supposing that Keith’s wife does not know about Keith’s evidence) what Keith says (‘I know that p’ or ‘I do not know that p’) will be compatible with Keith observing or trying to observe the general conversational rules. It follows that no substitutional implicature is involved in the bank cases.

14 See Fantl and McGrath (2009: 42).

something true, for it seems clear that we should not tell falsehoods.¹⁵ There might be some cases in which it's appropriate to tell a white lie, but it's hard to see why that would be the case in HIGH.¹⁶ Call this problem the 'inappropriate falsehood problem'.

A third problem points to the way in which one is willing to retract one's previous 'knowledge' assertion when the context changes. Suppose that, in LOW, Keith utters 'I know that the bank is open'. Suppose, however, that Keith's practical situation changes. Keith is now in HIGH, and he then utters 'I do not know that the bank is open'. Suppose someone challenges his new assertion by saying "But you said you knew that the bank is open!" A natural reaction for Keith is to retract his former assertion that he knows. But that is most unexpected if his assertion was true and appropriate. Further, we do not find similar retractions when it comes to other kinds of pragmatic implicatures.¹⁷ Call this problem the 'retraction problem'.

Fourth, standard pragmatic approaches appeal to pragmatic implicatures. One crucial feature of pragmatic implicatures is that they are cancelable.¹⁸ However, the assumed implicature in HIGH does not seem cancelable. For example, it seems that, in HIGH, an utterance by Keith of 'I know that the bank will be open but I cannot eliminate the possibility that it has changed its hours' or 'I know that the bank will be open but we should check' seems incoherent (see Cohen 1999: 60). Although advocates of standard pragmatic approaches have tried to explain why the supposed implicature seems difficult, if not impossible, to cancel, we typically do not find a similar phenomenon with other uncontroversial implicatures.¹⁹ For example, some writers have proposed to say that the implicature in question is nearly universal. But other

15 DeRose (2009: 114, 124).

16 See Fant and McGrath (2009: 41-42).

17 See Dimmock and Huvenes (2014: sect. 4).

18 See Grice (1989: 44), Levinson (1983: 114).

19 See Dimmock and Huvenes (2014). I should note that although I agree with Dimmock and Huvenes that there is a strong feeling of incoherence when the 'knowledge' ascription is made in the first person, it seems to me that this feeling disappears when it comes to the third person version. At least, it seems that an utterance of 'Keith knows that the bank is open, but he need to investigate further' sounds much less incoherent than an utterance of 'I know that the bank is open but I need to investigate further' (see also Stanley 2008: 40). In section 4 below, I explain the infelicity of the first-person version on the basis of the infelicity of uttering 'p but I should check whether p' and the fact that knowledge is factive.

nearly universal implicatures do not seem that difficult to cancel.²⁰ Call this problem the ‘cancelability problem’.

Finally, any proposed pragmatic approach must appeal to a general rule explaining how the implicature is calculated.²¹ Most accounts invoke the rule of relevance, but that seems unsatisfactory.²² For one thing, it appears that the semantic content of Keith’s ‘knowledge’ denial in HIGH *is* relevant. For if more than knowledge is required in HIGH, the fact that Keith does not know directly settles the issue. But, if so, that’s not the irrelevance of the semantic content which can lead the hearer toward the pragmatic content. The rule of relevance misses the mark. For another, assuming that knowledge is not always sufficient for action and that more than knowledge is required in HIGH, we still lack an explanation why Keith’s ‘knowledge’ ascription in HIGH is bound to generate the (false) implicature that Keith is in a *good enough* epistemic position to act, rather than the (true) implicature that he is *not in a good enough* epistemic position to act. In general, there is no reason to think that asserting that we satisfy epistemic standards weaker than the ones required in the situation is bound to communicate that we satisfy these higher epistemic standards. For example, suppose that knowledge that *p* is required in some context *C*. An utterance of ‘I’ve reasons to believe that *p*’ in *C*, literally expressing that we satisfy epistemic standards weaker than knowledge, does not generate an implicature that our epistemic position is *good enough* and satisfies the standards for knowledge. Rather, it precisely communicates that our epistemic position is *not good enough* and does not satisfy the standards for knowledge. This is entirely as we should expect given the rule of quantity (assert the stronger!). This observation shows that the approach under consideration still needs to provide an explanation why a ‘knowledge’ ascription in HIGH is bound to generate an implicature that an epistemic position stronger than knowledge is satisfied. Call this problem the ‘calculability problem’.²³

To sum up, standard pragmatic approaches face five pressing problems: the apparent truth-value problem, the inappropriate falsehood problem, the retraction problem, the cancelability

20 See Rysiew (2001: 496; 2007: 646) for this proposal and Dimmock and Huvenes (2014) for the reply. Brown (2006) notes that some implicatures are not easily cancelable. But, first, note that it is unclear that implicatures should be understood in pragmatic terms and, second, implicatures are not additive. If the supposed implicature in HIGH is additive, it is not similar to an implicature. See Dinges (2018).

21 See DeRose (2009: chap. 3).

22 Appeals to the rule of relevance are made in particular by Rysiew (2001; 2005; 2007), Brown (2006), Hazlett (2009).

23 See Dimmock and Huvenes (2014). Gerken (2017) proposes to explain the alleged pragmatic implicatures associated with ‘knowledge’ ascriptions and denials in terms of a heuristic associated with knowledge-based assessments. For criticisms, see Vollet (2018).

problem, and the calculability problem. These problems raise serious doubts as to whether standard pragmatic approaches can succeed. We might think, then, that psychological approaches are more promising. However, as I will now explain, psychological approaches face problems of their own.

3. Psychological approaches

In general, psychological approaches seek to explain the variability exhibited in cases like LOW and HIGH by focusing on whether the subject *believes* (in the sense relevant for knowledge) the target proposition. In this section, I consider the two main psychological approaches. I highlight the main problems they face, some already noted in the literature, some others new.

According to psychological approaches, Keith does not know in HIGH because, due to psychological factors, his belief (in the sense required for knowledge) is not formed. This should not strike us as particularly surprising, after all. In HIGH, Keith perceives a possibility of error and high stakes. It's an empirically well-established fact that perceiving a possibility of error and high stakes normally affects the formation of belief.²⁴ In most views of knowledge, knowledge requires a settled or outright belief. As a result, psychological approaches can easily grant and explain claims (1)-(4) above. If Keith does not believe the target proposition in HIGH (at least in the sense required for knowledge), Keith does not know in HIGH. Therefore, Keith's 'knowledge' denial is true and appropriate; Keith's 'knowledge' ascription would be false and inappropriate and Keith's 'knowledge' ascriptions and denial just seem to have the truth-value that they have in fact.

Regarding (5) and (6), it may at first seem that since psychological approaches grant that Keith knows that the bank is open in LOW, while he does not know that the bank is open in HIGH, they can accept knowledge norms for assertion and action.²⁵ Indeed, it's rational for Keith to act on and assert the target proposition in LOW, where he knows, but not in HIGH, where he does not know. As we will see, though, things are slightly more complicated.

In order to assess psychological approaches, it is useful to distinguish a simple and a sophisticated version. According to the simple version:

24 See in particular Hawthorne (2004: 164), Williamson (2005: 226; 234-235), Nagel (2008; 2010a, 2010b) and the references they quote.

25 It looks as though Nagel's approach tries to remain as neutral as possible on the issue of epistemic norms of assertion and action.

Psychological approach (simple) In HIGH, the subject's belief is not formed, but the subject could have formed a justified belief and had knowledge.

According to the sophisticated version:

Psychological approach (sophisticated) In HIGH, the subject's belief is not formed. If we assume, however, that his belief is formed in these circumstances, and given natural psychological assumptions about normal subjects, we should expect his belief to be ill-based, and therefore unjustified.

The simple version has been defended by Bach and the sophisticated version by Nagel.²⁶ I will consider their proposals in turn. As we will see, the major problem of these proposals is that they are insufficiently general. We can design cases in which they offer no adequate explanation. In addition, although more general, Nagel's approach faces further specific problems.

3.1 The simple psychological approach

Bach's view can be captured by the following claims:

1. High stakes related to *p* provide (normative practical) reasons to consider further possibilities of error and not to act as if *p*, i.e. not to close the issue.²⁷
2. Not closing the issue (tends to) raises the threshold for confident belief, and makes it more difficult for the subject to form or maintain a confident belief that *p*.²⁸

²⁶ Some authors take Nagel's view to be revisionary, thereby suggesting that, on her view, the (perceived) level of stakes is constitutively or conceptually related to the epistemic standards required by knowledge and epistemic justification (see, e.g., Sripada and Stanley 2012: §4). However, as made clear in **Psychological approach (sophisticated)**, and explained in further details below, Nagel's view does not add a necessary condition on knowledge or epistemic justification beyond those already recognized by conservative approaches. Her view is perfectly compatible with the conservative claim that if two subjects S1 and S2 have the same evidence *E* and believe that *p* on the basis of *E* *in the same way*, then S1 is justified or knows that *p* if and only if S2 is justified or knows that *p* (no matter the respective practical situations of S1 and S2). However, Nagel points out that, as a matter of fact, a normal person is not disposed to believe *in the same way* when aware that the stakes are high. Perceived stakes have a *causal* relevance for belief formation.

²⁷ Bach (2008: 83): "Since the size of the stakes increases the cost of being wrong, it also increases the range of possibilities one should guard against (...) One's *practical* reasons give one reason not to treat certain possibilities as closed, hence not to act as if they do not obtain" (my emphasis). See also Bach (2005: 77).

3. Therefore high stakes related to p influence the formation or retention of the belief and knowledge that p .²⁹

This approach has it that, in HIGH, Keith is somewhat mistaken about the strength or scope of his epistemic position. He implicitly thinks his evidence is insufficient to know the target proposition, and that his belief would be unjustified, whereas, actually, it would be justified. The evidence is only insufficient to act on the target proposition. In High, Keith displays “*excessive epistemic caution*” (Bach 2010: 117; My emphasis).

If Bach is right that in high stakes contexts beliefs are typically not formed because possibilities of error which are merely *practically* relevant (i.e. relevant for justified actions) are implicitly or automatically treated as *epistemically* relevant (i.e. relevant for justified beliefs), it should be possible on reflection to correct this mistake and to maintain a confident belief while considering further possibilities of error. At least, we could imagine people susceptible of maintaining their confident belief and their knowledge while (for practical reasons) they leave open certain possibilities of error. On this view, excluding the possibility of error is at most a psychological precondition (or an enabling condition) for meeting the threshold for confident belief. That’s not an essential condition for confident belief and knowledge.³⁰

28 Bach (2008: 83): “One’s practical interest explains the rise in the threshold of confident, settled belief, and thoughts of counterpossibilities make it more difficult for this threshold to be crossed” (my emphasis). See also (2005: 77): “One’s threshold for (confidently) believing a proposition is a matter of what one implicitly takes to be sufficient reason to believe it (...) Even if in fact one is in a position to know something, thinking one is not in a position to know it is enough to keep one from believing it (at least not without reservations) and to lead one, if it matters enough, to look into it further (...) However, it does not follow that the standards of knowledge go up (or that the threshold of justification is higher or that the range of relevant alternatives to be ruled out is wider).” See also Bach (2005, 30); (2010, 117-118).

29 Bach (2008: 83): “[T]he higher stakes raise the threshold of confident, doubt-free belief (...). Lack of doubt-free belief keeps one from meeting the doxastic condition on knowledge.” See also Bach (2005: 52, 77, 81).

30 According to a pragmatist interpretation of Bach, closing the issue is essential to believing confidently: there are normative practical reasons (not) to believe which explain why a settled belief is pragmatically unjustified in the high stakes situation under consideration. While this interpretation has some appeal, it does not fit well with other claims made by Bach. For example, Bach writes (2010: 118) “A different situation arises if you do have a settled and confident belief that p . Even then, given the cost of being wrong you may think that you need to make sure that p , by ruling out certain possibilities of error that ordinarily would be too remote even to consider, much less bother with.” It’s clear, then, that on Bach’s view it’s perfectly possible to have a settled belief and yet consider further possibilities of error. Bach (2008: 83) also writes that when the stakes are raised “One’s practical interest *explains* the rise in the threshold of confident, settled belief, and thoughts of counterpossibilities make it *more difficult* for this threshold to be crossed” (my emphasis). On his view, not closing the issue makes it more difficult to form a settled belief, but that does not thereby make that impossible. In addition, it can be granted that practical interests explain the rise in threshold without granting that this rise is pragmatically justified. To be sure, Bach (2010: 117) writes that the caution that one manifests in high stakes situations, and which can lead us to manifest excessive epistemic caution, “has a practical rationale and is therefore not irrational”. But that does not imply that a settled belief can be justified or unjustified by practical reasons. It would only if Bach also assumed that believing confidently necessarily required closing the issue. Be that as it may, the pragmatist interpretation does not undermine what I say below.

If so, Bach's proposal is obviously not friendly to knowledge or justified belief norms, for it allows possible high stakes cases in which the subject believes and knows that *p*, but without granting that in these situations it's rational for the subject to act on *p*. If some (epistemic) reasons can be sufficient to justify a confident belief that *p*, while they are still insufficient to justify closing the issue about *p* and acting on *p*, then knowledge, or justified belief, is not the epistemic norm of action.

I think Bach's approach is very plausible when it comes to dealing with HIGH. However, there are also third-person cases to consider, like Cohen's airport case:

AIRPORT. Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and responds "Yes I know it does stop in Chicago." It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, "How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute." Mary and John agree that Smith doesn't really know that the plane will stop in Chicago. They decide to check with the airline agent. (Cohen, 1999: 58)

In this case, Mary and John seem to appropriately deny 'knowledge'. They are in a high stakes situation, and they are assessing Smith, a low stakes subject who, they know, confidently believes the proposition on the basis of some evidence. Mary and John's intuitively appropriate (although arguably false) assertion that Smith does not know cannot be explained by appealing to the idea that Smith is seen as not confidently believing the target proposition. And if Smith's evidence is sufficient to justify his belief, Smith knows. How can Bach's approach explain this case?

The attributors are in a high stakes situation. If Bach's approach is correct, they do not confidently believe the target proposition. But when we do not confidently believe that *p*, we cannot coherently think that someone else knows that *p*. It is then incoherent to ascribe 'knowledge that *p*' to this person:

Belief condition on coherent 'knowledge' ascriptions (BCKA) If you do not believe that *p* (in the sense relevant for knowledge) it's incoherent to ascribe 'knowledge' that *p* to others.

BCKA is a very plausible claim, and it explains the intuitive appropriateness of Mary and John's (false) claim that Smith does not know.

However, what about cases in which both the attributor and the subject confidently believe the proposition? Modify AIRPORT and assume that Mary has double-checked:

AIRPORT – GOOD EVIDENCE. Mary and John have now checked with the airline agent. Mary's son, Bob, was with Mary and John when they overheard Smith's assertion. However, Bob didn't hear the ensuing discussion between Mary and John. Puzzled by the fact that they have double-checked with the airline agent, Bob asks: "But why did you ask the airline agent? Didn't you hear this passenger, Smith, saying that the flight stops in Chicago?" Mary replies "Yes but he didn't really know. The itinerary could have contained a misprint or they could have changed the schedule at the last minute."

Mary now has more information. She confidently believes that the flight stops in Chicago. Mary's assertion may seem appropriate (even if false) but BCKA cannot explain that. Still, Bach may appeal to the idea that Mary is biased or mistaken about the threshold required for epistemic justification. Her threshold for believing is still high (what she takes to be the epistemic standard for knowledge is still too high) and she wrongly takes Smith's evidence to be insufficient for knowledge.

Bach's approach provides a very plausible explanation for HIGH, AIRPORT, AIRPORT – GOOD EVIDENCE. However, crucially, it fails to account for a different type of case. Consider:

CONFIDENT HIGH STAKES SUBJECT. We have a paycheck to deposit at the bank. We arrive at the bank, and we see the long line. We overhear Keith saying that the bank was open on Saturday two weeks ago. We are in a low stakes situation and our belief threshold is normal. You tell me "If the bank was open on Saturday two weeks ago, it's open on Saturdays. It will be open tomorrow". I agree and, confidently believing that the bank will be open, we decide to come back tomorrow. Just before leaving, however, we learn about Keith's high stakes situation and we overhear that he replies to his wife, who has mentioned the possibility that the bank will be closed, "No. I know that the bank will be open." Puzzled by Keith's assertion, I tell you "He should not say that".

It seems to me that you should agree that Keith should not assert 'I know that the bank will be open tomorrow'. However, in the situation under consideration, it's clear that we (the assessors) and Keith (the attributor and the subject) confidently believe the target proposition

on the basis of evidence supposedly sufficient to know. If, in this situation, we judge that confident high stakes Keith should not say that he knows, that's not because we think that Keith does not confidently believe the proposition.

Although Bach's approach seems to be on the right track, it is insufficiently general. It cannot explain cases where it is stipulated that the assessor, the attributor and the subject all confidently believe the target proposition on the basis of what is supposed to be knowledge-level evidence and where, intuitively, the high stakes subject / attributor should not say 'I know that p'. As I will explain in section 4, we need a pragmatic approach to handle such cases. Before coming to that, however, it is useful to consider Nagel's more sophisticated approach.

3.2 The sophisticated psychological approach

According to Nagel's sophisticated approach:

Psychological approach (sophisticated) In HIGH, the subject's belief is not formed. If we assume, however, that his belief is formed in these circumstances, and given natural psychological assumptions about normal subjects, we should expect his belief to be ill-based, and therefore unjustified.

If Nagel's approach is correct, we avoid the problem raised by CONFIDENT HIGH STAKES SUBJECT. We may hold that the confident high stakes subject (/the attributor) does not know because his belief is unjustified. Therefore, we cannot truly ascribe 'knowledge' to this confident high stakes subject, and we can thus explain why we think that he shouldn't say 'I know'.

It's crucial to stress that the claim that the high stakes subject's confident belief is (expectably) not justified in high stakes cases is not based on the idea that the epistemic standards for knowledge shift with high stakes. Nagel's suggestion is rather that perceiving high stakes (expectably) compromises the possession of evidence, or at least the treatment of the (otherwise sufficient) available evidence.

In Nagel (2008), the proposal is couched in terms of "need-for-closure". Need-for-closure corresponds to the point where, as Kruglanski and Webster put it, "a belief crystallizes and turns from hesitant conjecture to a subjectively firm 'fact' ".³¹ The need-for-closure of an

31 Kruglanski and Webster (1996: 266), quoted by Nagel (2008: 287).

individual is determined by the integration of various motivational factors affecting his belief-forming processes to different degrees. Some factors, like pressure of time, cognitive costs, or perceived tediousness of the cognitive task, raise the need-for-closure. Others, like penalty for inaccuracy or pleasure taken in the cognitive activity, lower the need-for-closure. Some factors, like haste or distraction, are “accuracy-compromising”: they raise the probability of being biased. Given that knowledge or epistemic justification requires a well-founded belief, if a belief is perceived as based on an accuracy-compromising factor, the subject will be perceived as not knowing or not justified.³² In Nagel (2010), the proposal appeals to the notion of “epistemic anxiety”, where epistemic anxiety is a force or desire for increased cognitive activity (either in terms of greater amount of evidence or in terms of a more careful treatment of the evidence).³³ Likewise, epistemic anxiety can be overshadowed by accuracy-compromising factors in such a way that the subject will be perceived as not knowing or not justified. The relation between the notions of need-for-closure and epistemic anxiety is intricate. For our present purposes, it suffices to emphasize the crucial role that accuracy-compromising factors play in Nagel’s explanation of most cases.

To begin with, consider HIGH. Like in Bach’s proposal, Nagel’s approach has it that Keith does not know because he does not confidently believe the target proposition. The stakes are high. Keith feels strong epistemic anxiety; his need for closure is low.³⁴ His assertion ‘I do not know that the bank will be open’ is appropriate because it is true.

When it comes to cases like AIRPORT, Nagel’s approach departs from that of Bach in appealing to the hindsight bias, according to which we have a tendency to mistakenly project our own doxastic situation onto others.³⁵ Given this bias, Mary erroneously ascribes her low need-for-closure or the epistemic anxiety she feels to Smith and she thereby cannot see Smith’s confident belief as justified. Why? Smith’s evidence is relatively weak and, in Mary’s

32 Nagel (2008: 292).

33 Nagel (2010a: 414).

34 Nagel (2008: 289).

35 Nagel (2010a: 425). We should note that Nagel considers a slightly different case, from DeRose (2009): suppose that Lena, Thelma and Louise share the same evidence that John was at the office today (and they know that they share the same evidence). They saw his hat and they heard someone calling him. Now, Thelma at the tavern appropriately says that Lena knows that John was at the office. However, when the police ask for testimony that John was at the office, Louise appropriately says that she and Lena do not know that John was at the office. Nagel also proposes another possible explanation for the case involving Thelma, Lena and Louise. When the police ask Louise whether Lena knows, presumably, they are asking about Lena’s capacity to testify. However, if asked, given the epistemic anxiety she will feel, Lena will probably say that she does not know. That explains why it is appropriate for Louise to say that Lena does not know, thereby imparting that Lena could not testify that John was at the office.

point of view, Smith's confident belief must be the result of accuracy-compromising factors raising the need-for-closure or overshadowing his epistemic anxiety. This explains her (false) assertion that Smith does not know.

Consider now AIRPORT – GOOD EVIDENCE. Here too, we can appeal to the hindsight bias. If Mary projects her own doxastic situation onto Smith, in Mary's point of view, Smith should not have a belief unless he is under the influence of some accuracy-compromising factor. This explains Mary's (false) assertion that Smith does not know.³⁶

Finally, as anticipated above, contrary to Bach's account, Nagel's approach is in a position to offer a diagnosis for cases like CONFIDENT HIGH STAKES SUBJECT. We, the assessors, are aware that in this situation Keith is aware that the stakes are high. We then ascribe to Keith a relatively high level of epistemic anxiety or a low need-for-closure and we should not expect him to have a confident belief, unless he is under the influence of accuracy-compromising factors. If so, Keith does not know, which explains why we judge that it's inappropriate for him to say 'I know that the bank will be open'.

Since Nagel's account can explain the inappropriateness of Keith's assertion in CONFIDENT HIGH STAKES SUBJECT, it seems to fare better than Bach's proposal. Yet, as I will now explain, it is vulnerable to specific objections.

A first problem is the assumption that the presence of accuracy-compromising factors – that we supposedly need to postulate in order to make sense of the fact that Keith has a confident belief in CONFIDENT HIGH STAKES SUBJECT– is sufficient to explain why such a confident belief would appear unjustified. Indeed, explicitly mentioning an accuracy-compromising factor sometimes inclines us toward the intuition that the subject knows, or at least that he can appropriately say that he knows, rather than toward the opposite intuition.

To illustrate, consider a case from Schaffer (2006: 90) that Nagel uses to show that the need-for-closure (and thus the formation of belief) is affected by various kinds of practical factors:

HIGH AND FAST. On Friday afternoon, Sam is driving past the bank with his paycheck in his pocket. The lines are long. Sam would prefer to deposit his check before Monday, and indeed he has pressing financial obligations that require a deposit before Monday. His entire financial future is at stake. Sam remembers that the bank was open last Saturday, so he figures that the bank will be open this Saturday. He is right – the bank will be open. As Sam is about to stop to double-check the bank hours,

³⁶ Nagel (2010a: 426) also uses this strategy to explain our (alleged) intuition in so-called 'Ignorant High Stakes cases', in which the ascriber is aware that the subject faces a high stakes situation, but the latter is ignorant of that, and where it seems that the ascriber's 'knowledge' denial is appropriate (see Stanley 2005: 5). This case remains controversial, though (see Gerken 2017: 36).

he remembers that he promised to buy a present for his wife. She will be furious if he forgets – his whole relationship is at stake. The stores are about to close. Sam must choose. So Sam makes a split-second decision to drive past the bank and pick up a present for his wife instead, thinking that after all, the bank will be open this Saturday.

So, does Sam know that the bank will be open this Saturday?

Nagel (2008: 290) agrees that it seems more appropriate to say that Sam knows if he is short of time, and her explanation is that in that case Sam confidently believes because his need for closure is raised by pressure of time.³⁷ But time pressure is an accuracy-compromising factor in high stakes situations (at least typically). So if high stakes Keith's confident belief should be seen as unjustified why don't we see high stakes Sam (under pressure of time) as similarly unjustified?

Nagel might want to reply that pressure of time does not *always* compromise accuracy, and that it does not compromise accuracy in HIGH AND FAST. But we must then explain why, when it comes to explaining high stakes Keith's confident belief in CONFIDENT HIGH STAKES SUBJECT, we (allegedly) spontaneously and implicitly invoke an accuracy-compromising factor rather than a factor which does not compromise accuracy (like pressure of time in HIGH AND FAST). After all, postulating an accuracy-compromising factor is less charitable.

In brief, either we assume that pressure of time is an accuracy-compromising factor also in HIGH AND FAST, and we must grant that HIGH AND FAST shows that appealing to an accuracy-compromising factor is insufficient to explain why we judge that confident high stakes Keith does not know in CONFIDENT HIGH STAKES SUBJECT. For, in a similar situation involving Sam, we judge that Sam knows. Or else we assume that pressure of time is not an accuracy-compromising factor in HIGH AND FAST, but then it's unclear why appealing to an accuracy-compromising factor is necessary to explain the formation of Keith's confident belief in CONFIDENT HIGH STAKES SUBJECT.

This reveals a second problem for Nagel's approach. Appealing to accuracy-compromising factors is not the only possible way to account for the formation of confident beliefs in high stakes situations. Explanations involving other kinds of factors are possible, and they should be compatible with the ascription of a justified belief and knowledge. For example, we can say that confident high stakes Keith is psychologically abnormal and does not feel the epistemic anxiety normally related to the perception of high stakes. We can also say that

³⁷ See Nagel (2010a: 291).

Keith is practically irrational and does not care about the cost of error.³⁸ In such cases, his confident belief should be seen as justified.

To see the point, consider what Nagel writes about HIGH:

If we see the wife's request for additional evidence as rational, [Keith]'s refusal to take it seriously points to some epistemically problematic disposition in him. (Nagel, 2012, 682)

But, precisely, why would this refusal point toward an “epistemically” problematic disposition (concerning the justification of belief) rather than a “practically” problematic disposition (concerning the justification of action or assertion)? If Nagel's approach is meant to be non-sceptical and conservative, it has to grant that the wife's request is not epistemically rational, but merely practically rational. The possibility of error is not an alternative relevant to knowledge or justified belief.

We can also imagine a different case, in which high stakes Keith has a confident belief without being abnormal or practically irrational. Granted, empirical data suggest that people do not appeal to the same belief forming process (or method) when the stakes are low and when they are high. The standard view in psychology of reasoning, to which Nagel appeals, has it that people typically use two distinct systems or kinds of processes in reasoning.³⁹ But we can imagine a high stakes case in which the belief is carefully formed:⁴⁰

CAREFUL CONFIDENT HIGH STAKES SUBJECT. The same as HIGH, but Keith takes his time in assessing the evidence. He discusses with his wife, gives a weight to each clue and sums up the result in favour or against the belief that the bank will be open. At the end of the process, Keith acquires a confident belief that the bank will be open.

The method that Keith uses in this case is much more reliable than the method he is supposed to use in LOW (it involves type 2 processing).⁴¹ If the available evidence is sufficient to know, it should be sufficient for Keith to acquire a confident belief that the bank will be open.

³⁸Nagel grants that insensitivity to epistemic anxiety can sometimes be explained in terms of practical irrationality. For example, she writes: “the underlying reason why we should feel increased epistemic anxiety in high-stakes circumstances has to do with what is pragmatically desirable: the pragmatic benefits of increased accuracy offset the pragmatic costs of increased effort and time spent searching for evidence. If the (...) subject does not feel high anxiety in her current predicament, he is at risk of failing to maximize utility.” (Nagel, 2010a: 427)

³⁹ See Nagel (2010a: 409-412).

⁴⁰ See also Fantl and McGrath (2009: 44-46).

⁴¹ Nagel (2010a: 414) grants that there are two ways of improving our cognition: by acquiring further evidence or by thinking more thoroughly.

At least, if Keith acquires a confident belief, he cannot be seen as being influenced by accuracy-compromising factors (such as haste or distraction).

This point is emphasized by Fantl and McGrath (2009: 46), who correctly conclude that any conservative philosopher, Nagel included, has to accept the (in their view counterintuitive) view that in such a case Keith knows. It's clear, though, that Keith's assertion that he knows would seem inappropriate and that he should not act on the target proposition. According to Fantl and McGrath, that's because it intuitively still seems that Keith does not know, in spite of his careful belief formation.

Pace Fantl and McGrath, I side with Nagel (2008: 293; 2010: 422) in thinking that in this case – in which the subject's cognitive efforts are fully explained and where it's clearly postulated that there are no accuracy-compromising factors – the intuition that the subject does not know disappears (I will come back to Fantl and McGrath's objection in more detail in section 4). But, at least, it's incumbent on the conservative philosopher to offer a good explanation why it still seems inappropriate for Keith in such a case to say that he knows (and to act on the target proposition). Here, we might be tempted to appeal to standard pragmatic approaches, but we have seen that they are problematic. The challenge of explaining why confident and careful Keith should not say 'I know' remains.

Let us also note that an important consequence of granting that careful and confident high stakes Keith knows that the bank is open is that we must reject knowledge or justified belief norms of action and assertion, on pain of having to say that it's appropriate for Keith in such a case to act on the proposition that the bank will be open and to flat-out assert (that he knows) that that bank will be open.⁴² Nagel's approach seems much less friendly to knowledge or justified belief norms than we might have expected initially. (I will come back to the discussion of epistemic norms of assertion and action in section 4.)

Therefore, although Nagel accepts that, intuitively, the careful and confident high stakes subject knows, her approach is not well equipped to explain our intuition to the effect that he should not act on the target proposition and assert 'I know that the bank is open'.

In summary, the main problem of psychological approaches is that they are insufficiently general. Bach's approach provides no adequate explanation for cases like CONFIDENT

⁴² We might want to argue that careful and confident high stakes Keith is in a position to act on the relevant proposition without checking. However, given his epistemic position and the cost of error, the claim that careful and confident high stakes Keith should still check seems particularly difficult to deny (I know no philosopher denying that claim in print). A more promising strategy is to grant that careful and confident high stakes Keith should check while trying to maintain that that is compatible with knowledge being the norm of action. Williamson (2005) is an attempt of that sort invoking antiluminosity considerations. See Gao (2019) for criticisms.

HIGH STAKES SUBJECT. Nagel's approach proposes no explanation for CAREFUL CONFIDENT HIGH STAKES SUBJECT. In addition, I've raised two difficulties specific to Nagel's account. First, appealing to accuracy-compromising factors does not seem sufficient to explain why the belief under consideration should be seen as unjustified. Second, it's unclear why it's necessary to appeal to accuracy-compromising factors when we postulate that the high stakes subject has a confident belief.

4. Refined Invariantism

We have seen that non-sceptical and conservative approaches should grant that there are possible cases in which the high stakes subject has properly formed his belief and knows. Yet, in these cases, it still seems that the subject should not assert 'I know'. The psychological approach offers no explanation here. However, we have also seen that standard pragmatic approaches face important problems. In this section, I propose a new approach which supplements the psychological approach with a *non-standard* pragmatic approach appealing to the variability of the epistemic standards for appropriate assertion and action. I show that this combined approach is sufficiently general and that it avoids the problems faced by standard pragmatic approaches. To present the new approach I want to defend, I will start by a first-person case, namely, HIGH. I will then clarify how this approach also applies to third-person cases.

It has been stressed, against psychological approaches, that it is postulated in HIGH that "Keith remains as confident as before". Advocates of psychological approaches have replied, to my mind correctly, that if we postulate that Keith remains as confident as before (in the sense required for knowledge) that makes his assertion of 'I do not know that the bank is open' somewhat incoherent or inappropriate. As Bach (2005: 76) notes, when you take yourself as not knowing that p, in general, you are less confident that p than when you take yourself to know that p.⁴³

It follows that we must distinguish two high stakes cases which are mixed in HIGH. In the first one (HIGH 1), Keith appropriately asserts 'I do not know that the bank will be open'. In HIGH 1, we can explain why Keith's assertion is appropriate and seems true by granting that Keith does not remain as confident as before and no longer believes with sufficient confidence (that he knows) the proposition. The psychological approach provides an adequate explanation for this kind of case.

⁴³ See also Nagel (2010a: 421): "in announcing a decision to 'go in and make sure', HIGH certainly seems to be displaying lowered confidence, in some sense of that term."

In the second kind of case (HIGH 2), we postulate that Keith remains as confident as before, and hence, maintains his confident belief (and his knowledge). What would Keith say? As Fantl and McGrath (2009: 44) write, if Keith's confidence is still strong enough for Keith to be motivated to act, Keith would presumably say something like: "Aw come on, I know [the bank] will be open. We'll just come back tomorrow". However, they continue, "Intuitively, Keith's knowledge-claim, 'I know' is false", and, of course, it is irrational to come back without checking.

Let me note two things here. First, as said above, if it's clear that Keith has been careful in forming his belief, it does not intuitively seem to me that his claim to know is false (although I agree that it is inappropriate). Second, it's not clear that postulating that Keith maintains his confident belief amounts to seeing Keith as still motivated to act and willing to say "I know" and "We'll just come back tomorrow". After all, the same degree of confidence can be sufficient to move someone to act and assert in a low stakes situation and insufficient to move him to act and assert in a high stakes situation.

So, I think we should grant that it would be inappropriate for Keith to assert 'I know' and to act on the target proposition, whether or not Keith maintains his confident belief (that is, whether Keith is in HIGH 1 or in HIGH 2). In addition, we should accept, against standard pragmatic approaches, that if Keith has a confident belief (that he knows), it's inappropriate for him to say 'I do not know'.⁴⁴ Indeed, this is precisely what underlies the thought that there is some incoherence in postulating that Keith remains as confident as before in HIGH. For if Keith remains as confident as before (that he knows), why does he then deny that he knows? As explained in section 2, it's hard to see why a lie would be appropriate.⁴⁵ And if he sincerely denies that he knows, he is not as confident as before (that he knows), after all.⁴⁶ However,

44 See Fantl and McGrath (2009: 41).

45 According to Nagel (2010: 422), if in HIGH Keith is careful and confident it could still be "pragmatically rational for him to deny that he knows in order to placate his wife". However, a lie does not seem even pragmatically rational, for there is absolutely no need for Keith to lie to placate his wife. He just has to agree that he should check.

46 See Nagel (2010a: 421). It may seem surprising that we find Keith's knowledge denial intuitive given that it is postulated that Keith remains confident in HIGH. But note first that not everybody agree that the knowledge denial is intuitive. Some precisely think that it is at odd with the supposition that Keith remains as confident as before. Second, it's perfectly possible that those who intuitively see Keith's denial as appropriate implicitly appeal to accuracy-compromising factors to make sense of Keith's confidence, as Nagel suggests. Third, I suspect that when we assess the case, we easily overlook the condition that Keith remains as confident as before and that we directly focus on whether it is appropriate for Keith to deny that he knows. For, normally, when we face high stakes, our confidence is shaken. It's cognitively demanding and somewhat artificial to think of high stakes Keith as remaining as confident as before.

note that from the fact that we accept that confident high stakes Keith should not utter ‘I do not know’ in HIGH 2, it does not follow that we should accept that he can appropriately utter ‘I know’ in HIGH 2. Supposing that Keith justifiably maintains his confident belief, he could utter for example ‘I definitely think I know, but you’re right, we should make sure’, or ‘I really think the bank will be open, but you’re right we should check’.

The approach proposed here rejects the wrong dichotomy according to which Keith should either say that he knows or else that he does not know. This wrong dichotomy sets up the following fallacious challenge for conservative philosophers: either, if it is assumed that Keith remains confident and knows in HIGH (implicitly understood as HIGH 2), to explain how it could be appropriate for Keith to say that he does not know and why that seems true (standard pragmatic approaches try to meet this challenge); or, if it is assumed that Keith’s confidence is shaken in HIGH (implicitly understood as HIGH 1), to explain why Keith can’t know if the epistemic standards of knowledge do not shift (psychological approaches try to meet this challenge). But it’s clear that we can reject this false dichotomy. If Keith remains confident in HIGH (i.e. HIGH 2), the diagnosis seems to be that it’s inappropriate for him to say that he knows, but also inappropriate for him to say that he does not know.

In a nutshell:

Refined Invariantism applied to first-person cases. In HIGH, either Keith does not confidently believe that *p* (HIGH 1) or he does (HIGH 2). If he does not confidently believe that *p*, then it’s appropriate, because true, for him to say ‘I do not know that *p*’. If he confidently believes that *p*, then it’s not appropriate for him to say ‘I do not know that *p*’, and it’s not appropriate either for him to say ‘I know that *p*’. Keith could say something like ‘I think I know that *p*, but we should make sure’. This latter assertion is appropriate because it is true and does not pragmatically convey something false or inappropriate.

In rejecting the fallacious challenge, the conservative philosopher can divide and conquer. The first part of the challenge he now accepts is to explain why it *can be* appropriate (and why it *can* seem true) for Keith in HIGH to say that he does not know. This part of the challenge is adequately met by psychological approaches (which interpret HIGH as HIGH 1). The second part of the challenge is to explain why, if it’s postulated that Keith remains as confident in HIGH (which is to interpret HIGH as HIGH 2), Keith still should not say that he knows. The second part of the challenge can be met, I will now suggest, if we grant the independently plausible and well motivated claims that assertions and actions are governed by

variable epistemic standards and that satisfying knowledge-level standards is not always sufficient for appropriate assertions and actions.

The claim that the epistemic appropriateness of assertion or action is somewhat variable is virtually uncontroversial.⁴⁷ The claim that knowledge-level standards are not always sufficient for assertions and actions has been embraced and defended by many philosophers, and this claim is compatible with many different accounts of the norms of assertion and action.⁴⁸ In addition, it is crucial to note that the epistemic standards relevant for appropriate assertions and actions arguably do not all derive from the (supposed) epistemic norms governing assertions and actions. In light of this, even prominent advocates of knowledge norms have granted that the epistemic standards of appropriate actions and assertions are somewhat variable and can sometimes be stronger than knowledge-level standards.⁴⁹ Paradigmatic cases where these standards are higher include high stakes cases.

If these claims are correct, we have a simple explanation why, even if it is assumed that high stakes Keith has a confident belief and knows (HIGH 2), he should not assert (that he knows) that *p* and he should not act on *p*. Indeed, high stakes Keith does not satisfy the epistemic standards of appropriate assertion and action with regard to the proposition that (he knows that) the bank will be open.

The proposed explanation is not new. It invokes the warranted assertability manoeuvre according to which the unassertability of ‘I know *p*’ in HIGH (HIGH 2) is derived from the unassertability of ‘*p*’ and the factivity of knowledge. DeRose (2009: 90) considers explaining the bank cases in this way and calls this approach the ‘Generality Objection’ to contextualism.⁵⁰ Instead of putting the variability of the assertability in the specific truth-conditions of ‘knowledge’ ascriptions and denial, the Generality Objection places this variability in the general conditions for appropriate assertion. The reason why DeRose ultimately rejects the Generality Objection is that it cannot explain why it seems appropriate and true for high stakes Keith to say ‘I do not know that the bank is open’. As he puts it, “the

47 See DeRose (2009: 107).

48 See, among others, Brown (2008), Levin (2008), Reed (2010), Locke (2015), Gerken (2017).

49 For example, DeRose (2009) embraces (variable) knowledge norms but he also grants that secondary appropriateness requires a reasonable belief that one knows (2009: 93-94). Williamson (2005), who defends (invariant) knowledge norms, attempts to explain the apparent variability in terms of a general requirement to know (that one knows...) that one satisfies the relevant norm, where the relevant level of iteration is itself variable. See also Schulz (2018). For a discussion of the distinction between epistemic norms and epistemic standards, see Fassio (2017).

50 See also Pritchard (2010).

Generality Objection seems able to handle only our reluctance to claim knowledge, and seems ill suited to explain why we go so far as to deny that we know” (2009: 111). That’s why, in the end, we should accept that “there is something special about knowledge” (DeRose 2009: 110).

However, as explained above, we should not think that the Generality Objection fails if it cannot explain why Keith can appropriately deny that he knows. For if we suppose that Keith is in HIGH 1, that’s the psychological approach which is supposed to do the explanatory work. In contrast, if we suppose that Keith is in HIGH 2, it’s not appropriate for Keith to deny that he knows, and there is nothing to be explained that the Generality Objection fails to explain. Thus, in this framework, DeRose’s objection to the Generality Objection is undermined. Further, the invoked warranted assertability manoeuvre does not make the truth-conditions of ‘knoweldge’ ascriptions “special”. In this respect, Refined Invariantism seems better than Epistemic Contextualism.

In addition, it pays to note that the fact that it can be appropriate for confident high stakes Keith to say ‘I think I know’ suggests that revisionary approaches are not only unnecessary to handle the bank cases, but also dubious. For if it can be appropriate for Keith to say ‘I think I know’, that means that it’s possible that Keith knows in HIGH. But if revisionary accounts are right, it should be clear that Keith does not know. The belief that he knows should appear obviously false, and the assertion ‘I think I know’ inappropriate.

We have seen how first-person cases can be explained either by a psychological approach or by a non-standard pragmatic approach, depending on whether we postulate that the subject carefully believes (in the sense required by knowledge) the target proposition in the high stakes situation. Let’s now turn to the consideration of third-person cases.

Here, again, we can appeal to a psychological approach or to the non-standard pragmatic approach presented above, depending on the details of the case. We have seen that a psychological approach is well designed to explain third-person cases like AIRPORT and AIRPORT-GOOD EVIDENCE. In third-person cases of that sort the ascriber does not believe the target proposition or he believes the target proposition but he does not believe that the subject’s belief satisfies the epistemic standards for knowledge. The reason for that is that the ascriber is wrong, or biased, about the epistemic standards required for knowledge. Given this comprehensible mistake, it is appropriate for the ascriber to deny knowledge to the subject, although the subject actually knows. Refined Invariantism accepts the psychological explanation for such cases.

However, we have also seen that psychological approaches face a challenge when it comes to explaining third-person cases like CONFIDENT HIGH STAKES SUBJECT and CAREFUL CONFIDENT HIGH STAKES SUBJECT. In both of these cases, it is stipulated that the ascriber believes that the high stakes subject believes (in the sense required to know) the target proposition. It is also clearly stipulated in both cases that the ascriber is not biased and believes that the available evidence is good enough for knowledge. It is furthermore clearly stipulated in the latter case that the subject *correctly* bases his belief on the available evidence and that the ascriber does not think that the subject mistreats the evidence in question.

In these two cases, as explained in the previous section, it still seems that it would be inappropriate for the subject to overtly self-ascribe 'knowledge'. To explain that, Refined Invariantism appeals to the non-standard pragmatic approach: the epistemic standards of assertion are raised by the presence of high stakes, which explains why it would be inappropriate for the subject to assert (that he knows) the target proposition. And given that the subject is supposed to confidently believe (in the sense required to know) the target proposition, it also seems inappropriate for him to self-deny 'knowledge'. What we have suggested is that it would be appropriate for a subject in such circumstances to say "I think I know but I should make sure" or something along these lines.

Consider now the third-person ascriber. Given that the ascriber is aware that the subject confidently believes (in the sense required for knowledge) the target proposition on the basis of evidence sufficient to know and that the evidence is not mistreated, Refined Invariantism predicts that he could appropriately ascribe 'knowledge' to this subject. Crucially, that seems to be the right verdict. Again, following Nagel, if it is clearly stipulated that high stakes Keith believes on the basis of the available evidence and careful reflection that the bank will be open, the intuition that he does not know disappears and it's very unclear that it would be inappropriate to ascribe 'knowledge' to him. However, as explained above, that does not imply that the third-person ascriber should grant that Keith should not check or should assert that he knows. To repeat, a third-person ascriber could appropriately say "Keith knows, but he should make sure" or something like that.

In a nutshell:

Refined Invariantism applied to third-person cases. Either the ascriber does not believe the target proposition and/or (incorrectly) takes the subject not to believe the target proposition or to believe it in a wrong way -- for example on the basis of evidence insufficient to know (e.g. AIRPORT, AIRPORT- GOOD EVIDENCE) or in mistreating the evidence (e.g. CONFIDENT HIGH STAKES SUBJECT); or it is clear

to the ascriber that the subject believes the target proposition (in the sense required for knowledge) and does not mistreat the available evidence (e.g. CAREFUL CONFIDENT HIGH STAKES SUBJECT). In the first case, it is appropriate (although perhaps false) for the ascriber to deny ‘knowledge’ to the subject. In the second case, it is appropriate (and true) for the ascriber to ascribe ‘knowledge’ to the subject.

Thus, Refined Invariantism proposes to use a psychological approach to treat the third-person cases in which this approach does well, but it proposes to use a non-standard pragmatic approach to treat the third-person cases in which the psychological approach fails.⁵¹

To finish, let me highlight the virtues of Refined Invariantism. Refined Invariantism avoids the five pressing problems faced by standard pragmatic approaches. To see this, consider HIGH. First, if we interpret HIGH as HIGH 1, it is granted that Keith’s denial is true. That’s why what Keith says – ‘I do not know’ – seems appropriate and true (claims 1 and 3). In addition, Keith’s ‘knowledge’ ascription is false. That’s why it seems inappropriate and false for him to self-ascribe ‘knowledge’ (claims 2 and 4). At the same time, as noted, some (conservative) philosophers have the intuition that in HIGH Keith knows (they reject claims 3 and 4). Presumably, they interpret HIGH as HIGH 2. In such a case, Refined Invariantism grants that Keith’s ‘knowledge’ denial is false. We can thus easily explain why, in that case, as some philosophers have it, Keith’s ‘knowledge’ denial seems false and Keith’s ‘knowledge’ ascription seems true. And, of course it would be inappropriate for Keith to deny ‘knowledge’ in that case, for that’s false.⁵² It would also be inappropriate for Keith to ascribe ‘knowledge’, for the epistemic standards for appropriate assertion are too high in this context. As a result, the apparent truth-value problem is avoided, for there is no assumption here that true implicatures can make literally false assertions appear true or false implicatures can make literally true assertions appear false.

Second, Refined Invariantism avoids the inappropriate falsehood problem, for the proposed explanation does not need to assume that asserting a falsehood can be appropriate.

51 For example, consider the case of Louise being asked by the Police whether Lena knows that John was at the office (footnote 35). It is crucial to be explicit about whether Louise remains as confident as before. If Louise does not remain as confident, the psychological explanation is adequate. If Louise remains as confident and if it’s also clear that Louise’s confidence is not affected by a bias then it becomes unclear that she does not know and that she could appropriately say that she or Lena does not know. However, it still seems that Louise cannot appropriately assert that she or Lena knows. This is explained by the non-standard pragmatic approach.

52 Note that some revisionary philosophers agree that if Keith believes (that he knows) the proposition, he should not say ‘I do not know’ (e.g. Fantl and McGrath 2009: 41).

Third, we can also see the retraction problem in a different light. Suppose that in the new high stakes context under consideration, Keith's confidence is shaken. This makes the new context a case of HIGH 1. Keith's retraction is then easily explained by the psychological approach. Keith's threshold for believing is raised and, like John and Mary assessing Smith in AIRPORT, Keith (falsely) judges that he did not know. On the other hand, suppose that, in the new high stakes context, Keith maintains his confident belief. This makes the new context a case of HIGH 2. In that case, the retraction no longer seems appropriate.⁵³

Fourth, given that Refined Invariantism appeals to a warranted assertability manoeuvre invoking the epistemic standards for appropriate assertion and action, there is no cancellation problem. Indeed, there is no expectation that what is represented by the assertion – namely that the epistemic standards of assertion are satisfied – should be cancelable without incoherence. As is familiar, utterances of Moorean sentences – like 'p but I do not believe that p' or 'p but I do not know that p' – sound incoherent for the following pragmatic reason. By asserting 'p' you represent yourself as satisfying the relevant epistemic standards with respect to p for properly asserting 'p', and by asserting 'I do not believe that p' or 'I do not know that p', you represent yourself as not satisfying these standards. But it's incoherent to represent yourself as satisfying these standards and not satisfying them. Hence the uncancelability: either your assertion that p is inappropriate or your assertion that you do not believe or know that p is false and therefore inappropriate. Anyway, there is something inappropriate. In the framework developed here, similar considerations can be used to explain why an utterance of 'I know that the bank is open but I should check' by Keith in HIGH is bound to sound incoherent, even if it's true that Keith knows and should check.⁵⁴ Suppose that Keith does not satisfy the epistemic standards for appropriate assertion that (he knows that) the bank is open in this high stakes context, so that he should check before asserting. Therefore, it's inappropriate for him to assert '(I know that) the bank is open'. Suppose it's epistemically appropriate for him to assert this proposition. His assertion 'I should check' is false. Hence, the feeling of incoherence.

53 As Dimmock and Huvenes (2014) note, it should not be assumed that Keith's "only option is to retract his previous assertion. There might well be more than one appropriate response available to [Keith]. All that we are assuming is that if [Keith] were to retract his previous assertion and admit that he said something false, that would be a natural and appropriate response." I agree, for it's natural for Keith to lose confidence, and thereby knowledge, when he perceives that the stakes are high. In such a case, he cannot anymore ascribe 'knowledge' of the proposition to anyone. But we should not exclude cases in which Keith remains as confident as before, where a different response would be appropriate.

54 See Stanley (2008) for a similar manoeuvre regarding the incoherence involved in utterances of 'I know that p but I'm not certain that p / it's not certain that p'.

Last, it's clear that by invoking the variability of assertability conditions we are invoking a general phenomenon not specific to 'knowledge' ascriptions and denials. Refined Invariantism invokes a general rule, and it does not run the risk of being ad hoc.

Consider now the problems faced by psychological approaches. Their main problem is that they are insufficiently general. Bach's approach is insufficient to account for CONFIDENT HIGH STAKES SUBJECT and Nagel's approach (besides some other difficulties) is insufficient to account for CAREFUL CONFIDENT HIGH STAKES SUBJECT. In these cases, the high stakes subject is seen as having a confident belief based on evidence supposedly sufficient to know, and yet he should not say 'I know'. By invoking variable epistemic standards governing appropriate assertions, Refined Invariantism offers the needed explanation. In addition, Refined Invariantism is compatible with Nagel's contention that in some (or many) cases we see the high stakes confident subject's belief as unjustified because we see his belief as formed on the basis of accuracy-compromising factors. However, Refined Invariantism does not assume that this provides a necessary or sufficient explanation for what can happen in cases like the bank cases.

6. Conclusion

A certain number of cases presented in the literature highlight that our willingness to ascribe or deny 'knowledge' can shift with the stakes. While revisionary approaches have argued that that should be explained by the surprising assumption that the truth-value of these ascriptions and denials shifts with practical factors, conservative approaches have tried to invoke more traditional assumptions: that 'knowledge' ascriptions and denials have pragmatic implicatures, or that the perception of high stakes can affect the belief forming process. I have explained why these conservative approaches are ultimately unsatisfactory. I have proposed a third conservative approach instead – Refined Invariantism – which combines a psychological approach with a different pragmatic approach appealing to the epistemic standards of appropriate assertions and actions. I have shown why Refined Invariantism fares better than its rival conservative approaches.⁵⁵

55 A previous version of this paper was presented at the conference "Knowledge, Belief and Probability" in Hamburg in 2018. I would like to thank in particular M. Schulz, P. Rich, J. Koscholke, R Heil, F. Teroni and A. Logins for very helpful discussion. The work on this paper was supported by the SNSF research project 'Rationality and Reflection' (grant number 178039).

References

- Bach, K. (2005). The emperor's new 'knows'. In Preyer, G. and Peter, G., editors, *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, 51-89. Oxford University Press.
- Bach, K. (2008). Applying pragmatics to epistemology. *Philosophical Issues*, 18(1): 68-88.
- Bach, K. (2010). Knowledge in and out of context. In Campbell, J. K., Rourke, M. O., and Silverstein, H. S., editors, *Knowledge and Skepticism*, 105-36. MIT Press.
- Blome-Tillmann, M. (2014) *Knowledge and Presuppositions*. Oxford University Press.
- Brown, J. (2006). Contextualism and warranted assertability manoeuvres. *Philosophical Studies*, 130: 407-35.
- Brown, J. (2008). Subject-sensitive invariantism and the knowledge norm for practical reason. *Noûs*, 167-189.
- Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives*, 13, 57-89.
- Davis, W. A. (2007). Knowledge claims and context: Loose use. *Philosophical Studies*, 132(3), 395-438.
- DeRose, K. (1992). Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, 52, 913-929.
- DeRose, K. (2009). *The Case for Contextualism*. Oxford University Press.
- Dimmock, P., & Huvenes, T. (2014). Knowledge, conservatism, and pragmatics. *Synthese*, 191(14), 3239-3269.
- Dinges, A. (2018). Knowledge, intuition and implicature, *Synthese*, 195(6), 2821-2843.
- Engel, P. (2009). Pragmatic encroachment and epistemic value. In Haddock, A., Millar, A., and Pritchard, D., editors, *Epistemic Value*. Oxford University Press.
- Fantl, J. and McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford University Press.
- Fassio, D. (forth.). Moderate Skeptical Invariantism, *Erkenntnis*.

- Fassio, D. (2017). Is there an Epistemic Norm of Practical Reasoning? *Philosophical Studies* 174 (9): 2137-2166.
- Gao, J. (2019) Against the Iterated Knowledge Account of High-Stakes Cases. *Episteme*.
- Gerken, M. (2017). *On Folk Epistemology*. Oxford University Press.
- Grice, H. P. (1989). *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Hawthorne, J. (2004). *Knowledge and Lottery*. Oxford University Press.
- Hazlett, A. (2009). Knowledge and conversation. *Philosophy and Phenomenological Research*, 78(3), 591-620.
- Ichikawa, J. (2017). *Contextualising Knowledge: Epistemology and Semantics*. Oxford University Press.
- Kruglanski, A. W. & Webster, D. M. (1996). Motivated closing of the mind: “seizing” and “freezing.” *Psychological Review*, 103(2): 263-283.
- Levin, J. (2008). Assertion, practical reason, and pragmatic theories of knowledge. *Philosophy and Phenomenological Research*, 76(2), 359-384.
- Levinson, S. (1983). *Pragmatics*. Cambridge: Cambridge University Press.
- Locke, D. (2015). Practical certainty. *Philosophy and Phenomenological Research*, 90(1): 72-95.
- Lutz, M. (2014). The pragmatics of pragmatic encroachment. *Synthese*, 191(8), 1717–1740.
- MacFarlane, J. (2011). Relativism and Knowledge Attributions, in *Routledge Companion to Epistemology*, ed. Sven Bernecker and Duncan Pritchard. London: Routledge, 536-544
- Nagel, J. (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy*, 86(2), 279-294.
- Nagel, J. (2010a). Epistemic anxiety and adaptive invariantism. *Philosophical Perspectives*, 24(1), 407-435.
- Nagel, J. (2010b). Knowledge ascriptions and the psychological consequences of thinking about error. *Philosophical Quarterly* 60 (239): 286-306.
-

- Pritchard, D. (2010). Contextualism, scepticism and warranted assertibility manoeuvres. In J. K. Campbell, M. O'Rourke, & H. Silverstein (Eds.), *Knowledge and skepticism. Topics in Contemporary Philosophy*, Cambridge, MA: MIT Press, 85-103.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Noûs*, 35(4), 477–514.
- Rysiew, P. (2005). Contesting contextualism. *Grazer Philosophische Studien*, 69 (1):51-57.
- Rysiew, P. (2007). Speaking of knowing. *Noûs*, 41(4), 627-662
- Schulz, M. (2018) Decisions and Higher-Order Knowledge, *Noûs*, 51, 463-83.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.
- Stanley, J. (2008). Knowledge and certainty. *Philosophical Issues*, 18(1):35-57.
- Vollet, J.-H. (2018). The Warrant Account and the Prominence of 'Know'. *Logos and Episteme*, 9 (4): 467-483.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford University Press.
- Williamson, T. (2005). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *Philosophical Quarterly*, 55(219): 213-235.
-