

Against the iterated knowledge account of high-stakes cases

Jie Gao
KU Leuven

Abstract: One challenge for moderate invariantists is to explain why we tend to deny knowledge to subjects in high stakes when the target propositions seem to be inappropriate premises for practical reasoning. According to an account suggested by Williamson, our intuitive judgments are erroneous due to an alleged failure to acknowledge the distinction between first-order and higher-order knowledge: the high-stakes subject lacks the latter but possesses the former. In this paper, I provide three objections to Williamson's account: i) his account delivers counterintuitive verdicts about what it is appropriate for a high-stakes subject to do; ii) the high-stakes subject doesn't need iterated knowledge in order to be regarded as appropriately relying on the relevant proposition in practical reasoning; iii) Williamson's account doesn't provide a good explanation of why the high-stakes subject would be blameworthy if she were relying on the relevant proposition in her practical reasoning.

Keywords: practical reasoning; high stakes; higher-order reasoning; higher-order knowledge; Timothy Williamson

0. Introduction

Our ordinary epistemic assessments of action seem to suggest that knowledge sets the epistemic standard for appropriate practical reasoning. If I know that my car is parked in the yard, it seems that I am entitled to use that my car is parked in the yard as a premise in my practical reasoning. Conversely, if I fail to know that my car is parked in the yard, it seems that I am not entitled to the premise that my car is parked in the yard and hence would be vulnerable to blame if I used that premise in my reasoning. Such considerations motivate the following claim often referred as *the knowledge norm of practical reasoning*:

(KNP) One knows a proposition *iff* that proposition is an appropriate premise in practical reasoning.¹

Despite its popularity, KNP leads to an immediate objection to what is considered an orthodoxy in epistemology, *moderate invariantism*. Moderate invariantism can be characterised, roughly, as the view that epistemic terms such as ‘know’ invariably refers to an epistemic standard that i) exclusively depends on truth-relevant factors and ii) can be met quite easily. Following DeRose, the phrase ‘truth-relevant’ roughly denotes factors that “affect how likely it is that the belief is true, either from the point of view of the subject or from a more objective vantage point” (DeRose 2009: 24). Examples of truth-relevant factors are a belief’s safety and the reliability of its forming processes. By sticking to certain intuitive judgments, we can quickly derive a rejection of moderate invariantism. Suppose that for a low-stakes subject S it is not very important to be right about whether some proposition p is true, but for a high-stakes subject S^* it is very important to be right about whether p is true. Also suppose that S and S^* believe p on the same, good but not very robust, epistemic grounds (e.g., a memory about what happened two weeks ago).^{2,3} Intuitively, S can readily rely on p in practical reasoning, but S^* can appropriately rely on p only after taking some extra-precaution. Thus we have:

1. p is an appropriate premise for S ’s practical reasoning
- 1*. p is not an appropriate premise for S^* ’s practical reasoning

By applying KNP to S and S^* respectively, we get:

¹ Prominent upholders of the knowledge norm of practical reasoning include Hawthorne & Stanley (2008) and Williamson (2005) among others. Since the target of this paper is Williamson (2005), I use his formulation of the knowledge norm in that paper.

² Concrete examples include DeRose’s bank case and Cohen’s airport case among others (DeRose 1992: 912; Cohen 1999: 58).

³ Here I assume a sense of ‘epistemic grounds’ which includes only truth-relevant factors.

2. S knows p iff p is an appropriate premise for S 's practical reasoning
- 2*. S^* doesn't know p iff p is not an appropriate premise for S^* 's practical reasoning

From (1) and (2), we have:

3. S knows p .

And from (1*) and (2*), we get:

- 3*. S^* doesn't know p .

(3) and (3*) together are incompatible with moderate invariantism according to which if S knows p and S^* is as epistemically well-positioned with respect to p as S is (i.e., they share the same epistemic grounds), then S^* also knows p .⁴

In order to avoid a derivation of (3*), moderate invariantists have to reject either (1*) or (2*). On the assumption of the truth of intuitive judgments in (1) and (1*), most moderate invariantists have so far rejected KNP and hence (2*) (e.g. Brown 2008; Reed 2010; Gerken 2011). Although Timothy Williamson is also a moderate invariantist, as the main proponent of knowledge-first epistemology he has reason to endorse KNP. As he says, “without KNP the concept of knowledge would lose some of its significance: one reason why it matters whether you know something is that, if you do, you are entitled to use it in ways in which you would not otherwise be so entitled”

⁴ Another parallel argument against moderate invariantism with (1) and (1*) as premises can be built based on a meta-linguistic knowledge norm of practical reasoning. According to this norm, a first-person present-tense ascription of ‘know’ with respect to a proposition is true in a context iff that proposition is an appropriate premise for practical reasoning in that context (see Williamson 2005, pp. 227-8 for the full argument).

(Williamson 2005: 228). In his response to the above objection to moderate invariantism, Williamson (2005) thus suggests maintaining KNP and proposes an error theory of our intuitive judgment in (1*). According to this error theory, although p is an appropriate premise for both S and S^* , neither of them knows that she knows p , and thus neither knows that p is an appropriate premise. However, given the high stakes, S^* needs to know that p is an appropriate premise in order to be *regarded* as appropriately relying on p in practical reasoning. Mere knowledge of p is not sufficient (Williamson 2005). Since Williamson's account appeals to the lack of iterated knowledge in explaining away the intuitive judgment about high-stakes cases, we can name his account *the iterated knowledge account of high-stakes cases*.⁵

The aim of this paper is to provide three criticisms to Williamson's account. Here is the plan of the paper. In §1, I illustrate Williamson's account in more detail. In §2-§4, I provide three objections to Williamson's account: i) Williamson's account delivers counterintuitive verdicts about what it is appropriate for a subject to do in high stakes; ii) contrary to what Williamson claims, S^* doesn't need higher-order knowledge in order to be regarded as appropriately relying on p in practical reasoning; iii) Williamson's account doesn't provide a good explanation of why S^* would be blameworthy if she were relying on p in her practical reasoning.

1. The iterated knowledge account of high-stakes cases

⁵ Another challenge for moderate invariantists, derived from the same sort of cases, relies on a conflict between intuitive judgments about knowledge ascriptions: intuitively, it seems felicitous to ascribe knowledge that p to S but not to S^* . In response to this challenge, Williamson provides a psychological bias account. According to this account, it is natural for us to assign more weight to considerations telling against knowledge ascription when possibilities of error are made psychologically salient. Features such as the high practical costs of error for the subject or the ascriber, or the possibilities of error described in vivid and convincing detail can make the possibilities of error psychologically salient (Williamson 2005: 226). When we focus on considerations telling against the ascription of knowledge to the subject, we tend to withdraw the positive ascription of knowledge and to deny knowledge to the subject (*ibid.*: 234-5). Note that the psychological bias account cannot deal with the argument based on KNP considered in the main text. This is why Williamson introduces the account discussed in this paper.

Williamson's account of our assessments in (1) and (1*) has two key steps. First, by appealing to the non-luminosity of knowledge, Williamson argues that even though S and S^* know p , neither of them knows that she knows p . Second, Williamson argues that in order to be regarded as appropriately using p as a premise in practical reasoning, S^* (but not S) must possess *second-order knowledge*. Let us consider each step in more detail.

Concerning the first step, a condition is luminous just in case whenever one is in it, one is in a position to know that one is in it. According to Williamson, only trivial conditions are luminous; for instance, those that obtain in all cases or in none. By constructing a sorites series between a case in which the condition clearly obtains and one in which it clearly fails to obtain, Williamson argues that luminosity must fail close to the boundary between cases where the condition obtains and cases where it does not, just on the obtaining side (Williamson 2000: Ch.4). Neither knowing a proposition nor being an appropriate premise for practical reasoning are trivial conditions. Hence they are non-luminous conditions. It follows that in some cases one is not in a position to know that one knows q even if one knows q . Likewise, in some cases one is not in a position to know that the fact that q is an appropriate premise even if q is an appropriate premise. Given KNP, q is an appropriate premise iff one knows q . Hence, when q is an appropriate premise but one is not in a position to know that q is an appropriate premise, one in effect knows q without being in a position to know that one knows q (Williamson 2005: 230-1).

For what concerns S and S^* , due to the setting of the cases, according to Williamson, although both of them have the relevant first-order knowledge that p , neither of them is in a position to have second-order knowledge. This is because their knowledge falls into cases close to the boundary between knowledge and ignorance, just on the knowledge side. This seems to be plausible given the specific setting of the cases. All low-high-stakes case pairs are conceived in a way that the epistemic grounds shared by the low-stakes and the high-stakes subjects are not very strong, merely sufficient to convey the intuition that the subject in the low-stakes case knows.

As for the second step, according to Williamson, the fact that a subject relies on an appropriate premise without being in a position to know that it is appropriate provides some potential reason to question or criticise the decision. How harsh we should be with the subjects in such cases depends on how much is at stake. As Williamson says, “If not much, then it seems unreasonably pedantic to condemn the reasoning. But if matters of life and death are at stake, the charge that the agent was not in a position to know that the premise was appropriate becomes more serious” (*ibid.*: 230). Thus, given the practical situation of S^* (high stakes on whether p), in order to be regarded as appropriately treating p as a premise in her reasoning, it is not sufficient for S^* to be merely in a position in which it is appropriate for her to rely on p ; rather, S^* should also know that p is an appropriate premise.

Williamson seems to hint that the second-order knowledge requirement on high-stakes subjects is related to a corresponding requirement to engage in second-order reasoning about whether to trust the first-order practical reasoning. For example, he writes (where ‘Hi’ refers to S^* and ‘Lo’ to S):

[H]i has far more reason than Lo to check on such practical reasoning, to engage in second-order practical reasoning about whether to trust the first-order practical reasoning. Since Hi is in no position to know that the first-order premise that [p] is appropriate, the second-order premise that the first-order premise is appropriate is, although true, inappropriate (given [KNP]). Thus second-order reasoning is in no position to give a clean bill of health to first-order reasoning based on the premise that [p] (*ibid.*: 232-3).

Given the high stakes situation that S^* faces, S^* has far more reason than S to check whether she can rely on her first-order reasoning, *i.e.* whether she can trust that p is an appropriate premise for

her practical reasoning and thus, given KNP, whether she knows p . Williamson suggests that this second-order reasoning requirement is prudential: it would be highly imprudent for S^* to directly rely on p without engaging in a second-order reasoning about whether it is appropriate to use p as a premise in practical reasoning, and without reaching a positive answer to that question through this second-order reasoning (*ibid.*: 233).

Combining the non-luminous condition with the second-order knowledge requirement, Williamson explains the intuitive judgment in (1*). According to the second-order knowledge requirement, S^* needs to have second-order knowledge of p in order to be regarded as appropriately using p as a premise in her practical reasoning. But according to the non-luminosity condition, S^* is not in a position to have second-order knowledge of p . Thus we deem that all things considered it is inappropriate for S^* to use p as a premise in her practical reasoning, even though p is indeed an appropriate premise for her practical reasoning.

Williamson also adds that if the stakes are high enough, a prudent subject should engage in even higher orders of reasoning—third, fourth, and so on—about whether to trust the previous-order reasoning (*ibid.*: 233). In addition, Williamson shows how a failure to have $n+1$ iterations of knowledge that q in deliberation could end up leading to a self-denial of knowledge that q . Williamson invites us to consider a dialogue in which one interlocutor first asks another (who could also be herself) whether q is the case; then, provided a positive answer, the interlocutor asks whether she can provide warrant for the answer she just gave. The interlocutor continues asking the same question for each positive answer. Sooner or later, the interlocutor under interrogation would run out of warrant. Williamson argues that when this happens, previous positive answers will in turn be destabilised in a domino effect. Similar consequences apply when one considers whether one has warrant for various levels of higher-order knowledge. When one finds out that she lacks warrant for some higher-order knowledge of q , all lower-order knowledge of q will be in jeopardy as well. This shows that a failure of some higher-order reasoning in providing justification

for the lower-order reasoning would in the end hinder one from relying on the target proposition in her first-order practical reasoning (*ibid.*: 233-4). This allows moderate invariantists to deal with cases in which high-stakes subjects have second-order (or even higher-order) knowledge. In such cases, the seeming lack of first-order knowledge and of warrant to rely on it can be accounted in terms of a lack of some higher-order knowledge.

2. Counterintuitive appropriateness

In this section I argue that Williamson's account delivers very counterintuitive verdicts about what it is appropriate for a subject to do in high stakes. According to Williamson, both S and S^* satisfy KNP's conditions for *appropriate* use of p as a premise in practical reasoning. Two clarifications are in order about the notion of 'appropriateness' used in KNP. First, there is an issue about how to understand KNP as an epistemic norm. According to one obvious and widely acknowledged understanding, KNP is an epistemic norm in the sense that it demands that a certain epistemic condition with respect to q be satisfied in order to rely on q in practical reasoning. In a narrower sense, KNP is an epistemic norm in the sense that its normative source comes from an epistemic standard as opposed to a different normative standard (prudential, moral, aesthetic, *etc.*).⁶ It's likely that Williamson has the narrower sense in mind.⁷ Here, for the sake of argument, I will assume this interpretation. Second, following Hawthorne & Stanley (2008: 578), 'appropriateness' in KNP should be conceived as permissibility rather than obligation.

With these clarifications in mind, it doesn't sound quite right to say that S^* would do anything epistemically permissible if she were using p as a premise in practical reasoning. We might be able to see the point more clearly by looking at a concrete example. Since Williamson uses Cohen's airport case as the target example in his paper, the case I suggest is a modification of that one:

⁶ On this point see, for example, Crisp (2005). For a discussion of different sources of normativity see, for example, Broome (2013), pp. 26-7 and Ch.7.

⁷ Compare to his discussion of the norm of assertion in Williamson (2000), Ch.11.

Trustful Airport

Mary and John are at the Los Angeles 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 another 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.” Mary and John have to deliver an organ for an urgent transplant on a patient in Chicago. They are aware of the fact that in some rare cases the itinerary could contain a misprint or the schedule could have been changed at the last minute. Still, they rely on the information they overheard from Smith without any further check about the itinerary. It turns out that the itinerary used by Smith is reliable and provides the correct information.

It seems that Mary and John should collect more evidence and should have a stronger epistemic position in order to be justified enough to rely on that information in their practical reasoning. But according to moderate invariantism, since Smith knows that the plane stops at Chicago (henceforth r), and Mary and John acquire a true belief of r based on Smith’s reliable testimony, Mary and John also know r . Then according to Williamson’s account, it is appropriate for Mary and John to use r in their practical reasoning (assuming that they acquire knowledge by that testimony).

If we take seriously Williamson’s account, we should be able to distinguish at least two kinds of evaluation about the subjects’ practical reasoning. One evaluation would be about the epistemic permissibility of relying on r in one’s practical reasoning; the other would concern other evaluative standards relevant in judging the subject’s decision-making, *e.g.* prudence. In terms of the habit of decision-making exhibited, it is indubitable that Mary and John are utterly imprudent in relying on r without searching for any further evidence for r .

Now, when there are multiple evaluative standards according to which an action can be assessed and they deliver opposite verdicts, normally we can easily tell those standards apart from one another and acknowledge a conflict between the respective evaluative judgments (at least from a third-person perspective fully informed about the facts). For instance, we can easily distinguish epistemic assessments from moral assessments in the following case concerning assessments relative to assertion: While I may know that the fugitive is in the basement, and so satisfy the epistemic standard required to appropriately assert that the fugitive is in the basement, this assertion would violate a moral rule if my behaviour hinted at the presence of the fugitive to the enemy soldiers (McKenna 2015: 4).

Consider another case sharing the same structure of evaluation as the case of Trustful Airport. Suppose that Jimmy stops at the red light of a very busy crossroad. Suppose that he is in a country in which people respect traffic regulation only ‘moderately’. So he knows that without paying careful attention to vehicles passing by, he could end up having an accident. Nonetheless when the traffic light turns to green Jimmy immediately starts the engine and drives off without checking any further. Luckily, Jimmy passes the crossroad without incurring any accident. In this case, we can easily distinguish two levels of assessment, one positive and one negative: according to the traffic regulations, Jimmy hasn’t done anything wrong and should not be subject to any sanction considered by the law. On the other hand, Jimmy’s action was imprudent. He should have checked more carefully whether some vehicle was coming from the other direction—or, as Williamson may suggest, he should have engaged in second-order reasoning about whether it was appropriate for him to proceed.

However, in Trustful Airport, a similar distinction between different standards seem to be absent. Intuitively we (as third-person observers fully informed of the facts) don’t hold Mary and John as epistemically appropriate but prudentially irresponsible. Rather, it seems that our assessments about their decision to rely on the information they overheard are completely negative.

In this case, we can only recognise a unique negative evaluation (be it prudential or epistemic), not two evaluations driving different directions as it would be the case if Williamson were right. Thus holding, as Williamson does, that Mary and John have met the knowledge norm of practical reasoning is at most a theoretical speculation not supported by any intuitive judgment.

One might challenge this argument by asking why normative standards should fit with our intuitive judgments. In particular, one may argue that intuitive judgments go wrong in such cases precisely because these are borderline cases in which the subject is blind about whether she knows or not, as Williamson's account predicts. However, this possible reply misses the point. First, even if the subject in a borderline case were not in the position to recognize the double evaluation, we as third-person assessors fully informed about the facts and the epistemic position of the subject should have no problem in discerning these evaluations. Second, even from a first-person perspective, we should distinguish between on the one hand knowing that there is a norm and what it requires in general, and on the other hand knowing whether the conditions in the current circumstances conform with the norm's demands. For example, consider someone driving through a crossroad who knows that the law requires stopping at red lights but is not in the position to discern whether the light is green or not due to scarce visibility conditions. In this case, one knows that there is a norm requiring from her a specific thing but cannot know whether her action fulfils that requirement. Similarly, Williamson's account predicts that the high-stakes subject does not know whether she knows q and whether it is appropriate for her to rely on q . However, it does not predict that the subject is blind about whether the epistemic and prudential standards are in force in the context. But our intuitive judgment about relevant cases is not merely that the subject is uncertain about the assessments according to these standards. Rather, in these cases, intuitively there is only one standard in force involving a negative assessment on relying on q , also from the subject's perspective. In other words, while Williamson's account predicts that the *assessments* relative to these norms (whether the subject is acting appropriately in the circumstance) may not be

transparent to the subject, it does not say anything about whether the *enforcement* of the norms is transparent to her (whether the subject knows, for example, that there is a knowledge norm governing the use of propositions as premises in her reasoning). But in order to avoid my objection, one would need the latter type of blindness, not the former.

3. Higher-order reasoning/knowledge and practical rationality

The second objection focuses on the higher-order reasoning/knowledge requirement in Williamson's account. More precisely, I will question the following claim:

(HORK) Engaging in higher-order reasoning and/or having iterations of knowledge of q is necessary and sufficient in order for a high-stakes subject to be regarded as appropriately using q as a premise in her practical reasoning.

If HORK is false, Williamson's account of our intuitive judgments about the problematic cases in terms of second-order reasoning/knowledge is undermined.

My objection to HORK relies on a preliminary clarification of the relation between higher-order reasoning/knowledge and degrees of first-order warrant. It's not clear how we should understand this relation. According to one understanding, having higher-order knowledge of q is related to the strength of epistemic position with respect to q , measurable in terms of degrees of first-order warrant. For instance, one may argue that the closer to epistemic certainty the first-order knowledge is, the safer the corresponding higher-order beliefs are, and the higher the order of knowledge the subject is in a position to have. Since under such an understanding, having further iterations of knowledge is made possible by having stronger first-order warrant, one may question the relevance of talking of higher-order knowledge. Indeed, first-order warrant may be doing all the justificatory work, and an account of the relevant cases that focuses directly on

degrees of first-order warrant (e.g. Gerken 2011, 2015, forthcoming) would be a simpler, straightforward alternative.

In fact, I think that higher-order knowledge should be carefully distinguished from degrees of first-order warrant. First-order warrant and knowledge iterations are very different properties. This is evident if one observes that sometimes certain methods are good for the acquisition of one of these properties but not the other (e.g., higher-order reasoning can increase the number of knowledge iterations without adding first-order warrant). Higher-order knowledge is about lower-order attitudes, not directly about the strength of one's epistemic position with respect to the target proposition. Performing second-order reasoning and consequent acquisition of second-order knowledge about q cannot alone strengthen the first-order warrant with respect to q . Rather, first-order warrant seems to depend exclusively on first-order evidence supporting q .

It is worth noting that the above considerations stand even if one conceives first-order warrant in terms of reliability or safety of one's first-order belief. Williamson observes that belief's safety from error only grants that one is *in a position to* have a certain number of knowledge iterations (see Williamson 2000: 116 and Appendix II). Nonetheless, for actually possessing this iterated knowledge, the subject should draw deductively the pertinent conclusions from what she knows. This implies that when the safety of first-order belief necessary for one's being in a position to have these knowledge iterations is already in place, higher-order inferential reasoning can increase the number of knowledge iterations without adding reliability to first-order belief and first-order warrant.

Moreover, to the extent that degrees of first-order warrant are much more fine-grained than the number of iterations of knowledge, it follows that the two properties have different extensions. Assume, as is plausible in many ordinary cases, that degrees of first-order warrant are related to degrees of safety of first-order belief: the stronger the warrant for believing q , the more remote the cases in which q is false from those in which q is true, and the safer the belief that q . Now, in the

Williamsonian framework, different degrees of safety of first-order belief can grant the same safe margin from error sufficient for, and only for, putting one in a position to have a specific number of knowledge iterations K^n . This implies that a specific number of knowledge iterations is compatible with different degrees of first-order warrant. On a topological conception of safety like that discussed by Williamson (2000: §5.3), we can think of margins for error as limits of regions in an n -dimensional Euclidean space. The distance between the margin for safely believing q and the margin for safely believing q can occupy several points in the region, each of which corresponds to a different degree of reliability or first-order warrant. This implies that the reliability and first-order warrant of one's belief can be strengthened without necessarily enabling the subject to be in a position to have more iterations of knowledge.⁸

Furthermore, there are also reasons to separate the strength of first-order warrant from first-order belief's reliability (and consequently from the knowledge iterations one is in a position to have). Consider a trivial condition C . Examples include conditions one is always in (e.g. I exist), and conditions which cannot be unsafely believed because they obtain in every possible world (e.g. necessary truths and tautologies). According to Williamson, C is a luminous condition, one that a subject is in a position to know that it obtains whenever it obtains. Consider a belief about condition C . Such a belief enjoys maximal safety. Williamson's anti-luminosity argument cannot prevent possible infinite knowledge iterations about this condition. However, it seems that no matter how good the subject's evidence in support of the obtaining of C already is, new evidence would further strengthen her first-order warrant for that condition. For example, no matter how

⁸ If one finds a topological conception too abstract, one is free to think of more concrete examples in which new evidence is sufficient to strengthen one's belief reliability but not enough to pass the threshold for making the belief reliably reliable. For example, I can discriminate a slight increase in the heat of a surface between time t_1 and t_2 . Furthermore, my perception of the heat at both times is of a degree sufficient to grant the safety of my belief that from t_1 the surface has been more than 30 degrees Celsius, but insufficient to grant a safely safe belief in the same proposition. Thus, the reliability of my belief at t_2 is higher than the reliability of my belief at t_1 , but both of them are not enough to make my belief reliably reliable. See Williamson (2000: §5.3) for further discussions and examples.

well supported my maximally safe belief in the law of excluded middle is,⁹ if I come to know that some famous mathematician discovered a new theorem which further confirms the truth of the law, my first-order warrant for my belief will be further strengthened by this new evidence.¹⁰

It follows from the above discussion that there are cases in which a subject is in a position to acquire infinite knowledge iterations but doesn't have maximal first-order warrant. We can also conceive opposite cases in which a subject has very robust first-order warrant for q but does not have higher-order knowledge of q . We have already appreciated the fact that for actually possessing higher-order knowledge, the subject should deduce the pertinent conclusions from what she knows. One can have a very strong first-order warrant but fail to have higher-order knowledge simply because one doesn't engage in deductive second-order reasoning and thereby doesn't form the relevant higher-order beliefs.¹¹ In addition, there are also cases in which it is at least physically impossible for the subject in possession of a strong first-order warranted belief that q to form higher-order knowledge that q . Such cases are possible with subjects who have difficulties in forming higher-order attitudes and engaging in higher-order reasoning, such as young children and subjects who lack the concept of knowledge.¹²

Once we accept that having higher-order knowledge can diverge in important ways from having strong first-order warrant, we are in a better position to assess HORK. First, engaging in higher-order reasoning and having multiple iterations of knowledge seem to be unnecessary for being

⁹ I am here assuming that this truth is presented to me in a simple tautological guise, i.e., I have the concepts to formulate the tautology. See Williamson (2000: 107-108).

¹⁰ The latter consideration relies on the crucial point that while safety from error is a modal property of belief, warrant is a property related to the actual support that a certain body of evidence provides to a proposition. Some Bayesian epistemologists have suggested specific methods to measure this kind of support. See Joyce (2005) for an overview.

¹¹ This is easy to imagine if we consider cases in which the only way in which a subject can acquire higher-order beliefs is through higher-order reasoning, but the subject didn't yet engage in any such higher-order reasoning.

¹² Williamson recognizes this type of case, see Williamson (2000: 95, 107-108, 115). We can also conceive cases in which it is metaphysically impossible for the subject in possession of a strong first-order-warranted belief that q to form higher-order knowledge that q . Consider the case in which q is the proposition that the subject has no second-order attitudes.

regarded as appropriately using q as a premise in practical reasoning in high stakes situations in which the subject's epistemic position with respect to q is extremely strong or reaches the maximal degree. Consider an analogue of the bank case in which the evidence held by Keith that the bank will be open tomorrow (henceforth b) is that just a minute ago he read the bank's opening hours in front of the bank and also got a confirmation about that from the staff working at the reception desk.¹³ In this case, it seems that the strength of Keith's epistemic position with respect to b is pretty robust. Given the massive evidence Keith has for the proposition on which he acted, should we really blame him for not asking himself whether he really knows b , whether he knows that he knows b ... and for not forming a second, third... n -order belief about the matter? There is a clear intuition here that Keith is fully blameless in this situation. It seems obvious that Keith cannot be criticized for directly relying on b in his practical reasoning without first engaging in higher-order reasoning and forming higher-order knowledge of b . His first-order warrant is more than sufficient to justify him to rely on b .

Assuming HORK has other implausible consequences. First, according to Williamson's account, a progressive increase of stakes would give rise to an increasing demand of higher-order reasoning and knowledge iterations in order for the subject to be regarded as appropriately relying on a proposition for practical reasoning. Now, while it still seems somewhat intuitive that in some high-stakes contexts one should engage in second-order reasoning¹⁴, it sounds rather odd to also demand the subject to engage in a third-order reasoning about whether to trust the second-order reasoning about whether to trust the first-order reasoning. After all, the decision to be made in such cases is about the first-order question whether to treat q as a premise in practical reasoning. A third-order reasoning seems to be simply off the track in adjudicating the initial question

¹³ I conceive the case as one in which the source of information is more reliable than the original bank case. If one doesn't find the case helpful, one is free to change the case in accordance with what she thinks is necessary in order to select the reliable source.

¹⁴ After all, in such circumstances, it is reasonable for one to be reflectively conscious of one's epistemic position with respect to q when the practical costs for being wrong about q are severe, and a failure to engage in a second-order evaluation would be reckless and irresponsible.

concerning what to do. In this respect, appealing to higher-order reasoning in accounting for intuitive judgments about paradigmatic high-stakes cases appears to be a misplaced strategy.

It is also excessively pretentious to ask high-stakes subjects to engage in many levels of reasoning and possess many iterations of knowledge. For one thing, when ordinary people engage in deliberation about whether to perform a certain action, they rarely or never raise questions about whether they know the premises in their deliberations, let alone questions about even higher-order knowledge (cf. Greco 2014: 170). Concerning the maximal iteration of knowledge one can properly attribute, there is empirical evidence that neurotypical adults can only reliably track higher-order *interpersonal* mental state attribution, as in “Mary thinks that Tom doubts that Vivian hopes to stay”, at most for five levels (Kinderman *et al.* 1998; Stiller & Dunbar 2007). Presumably, higher-order *intrapersonal* knowledge attributions are subject to similar or even stricter natural limits. Engaging in very high orders of reasoning about knowledge iterations is definitely impossible for normal human beings given our inability to properly track many levels of reasoning.

Having higher-order knowledge and engaging in higher-order reasoning may not be necessary for one to be regarded as appropriately using q as a premise in practical reasoning in high stakes, but would it be at least sufficient? There are reasons to think that the answer is no. One may say that higher-order knowledge of q is indirectly relevant to q . But still merely having higher-order knowledge doesn't necessarily strengthen the warrant for q . As we saw above, there are trivial and hence luminous and maximally safe conditions for which the anti-luminosity argument cannot prevent infinite knowledge iterations. Now, suppose a subject $S^\#$ has several knowledge iterations of q about a luminous condition C . No matter how many knowledge iterations $S^\#$ has, we can conceive cases in which stakes on q are so high that $S^\#$ would not be regarded as appropriately using q as a premise in practical reasoning. Many share the intuition that it would be inappropriate

and imprudent to bet on any proposition whatsoever when stakes on q are sufficiently high, including propositions we take to be tautologies.¹⁵

There are cases in which engaging in higher-order reasoning and acquiring higher-order knowledge would even be regarded as inappropriate and detrimental for subjects in high stakes. These cases constitute counterexamples to both necessity and sufficiency directions in HORK. Consider, for example, cases in which an exclusive focus on first-order considerations is mandatory in order to complete important tasks. Take a case in which a pilot is trying to land a full-loaded aircraft on a river with broken engines. Any reflection about first-order practical reasoning (e.g., whether she really knows that pushing further a lever would have such and such consequences) could distract her from completing crucial procedures for a safe landing.¹⁶ Similar considerations apply to higher-order knowledge: imagine the pilot has not yet formed a higher-order belief about what she knows concerning emergency landings and is in the position to form it only by engaging in higher-order reasoning. Since she shouldn't engage in higher-order reasoning, she may also lack the corresponding second-order knowledge.

In sum, in this subsection I have argued that engaging in higher-order reasoning and/or having higher-order knowledge seems neither necessary nor sufficient for the subject to be regarded as appropriately relying on the target proposition in practical reasoning in certain high stakes situations. This constitutes a serious problem for Williamson's account. As shown in this section, what really matters in determining whether it is appropriate to use q as a premise in one's practical reasoning in a context is the degree of first-order warrant for that proposition, and these are not obviously related to the number of knowledge iterations.

¹⁵ See e.g., Hawthorne (2004: 29, fn. 72), Hawthorne and Stanley (2008: 587), Fantl and McGrath (2009: 189), Reed (2010: 228-29).

¹⁶ The case is inspired by a similar one in Markovits (2011: 157) which in turn draws on a real case. In the same article you can find other similar cases.

4. Non-luminosity, higher-order belief and blameworthiness

For the sake of argument, let us grant with Williamson that it is reasonable to regard S^* 's failure to engage in a second-order reasoning as imprudent and that having or lacking higher-order knowledge plays an important role in explaining why we would blame S^* for relying on p in her practical reasoning. In this section, I want to question that a lack of second-order knowledge *for the reasons provided by Williamson's anti-luminosity argument* can do this explanatory work. For Williamson, the lack of second-order knowledge is due to the failure of a safety requirement on belief. This cannot explain why S^* should be considered imprudent for failing to have second-order knowledge. In general, blaming someone for being imprudent requires some wrongdoing for which the subject is fully responsible, and thus which is recognisable from her own perspective. A proper explanation of this blame in the second-order knowledge failure would require that this failure be due to a lack of reasonable second-order belief—a belief that one knows q . Unfortunately, Williamson's anti-luminosity argument doesn't show failure of second-order knowledge due to failure of reasonable second-order belief.

Let me elaborate this point a little further. In the anti-luminosity argument, Williamson aims to show that any non-trivial condition is not luminous. The original argument considers cases involving indiscriminable transitions between a situation in which a condition clearly obtains (e.g., feeling cold) and one in which it clearly fails to obtain (i.e. feeling hot). In these cases there is a point at which the subject is confident that, for example, she feels cold at a time α , which is true but unsafe, *i.e.* such that at a very close time α_{i+1} the subject is still confident that she feels cold but it is not true that she feels cold. Luminosity fails precisely in these circumstances, in which one's confidence about a relevant proposition (e.g., that one feels cold), although true, is not safe from error, *i.e.* it is not reliably based. The argument applies to every other non-trivial condition, including knowledge itself.

This specific feature of the argument is particularly important for our discussion of Williamson's account: if S^* 's knowledge that p is non-luminous due to the reason provided by the anti-luminosity argument (i.e., lack of safety), it is compatible with the conclusion of the anti-luminosity argument that S^* is confident enough to believe that she knows p . Indeed, Williamson's argument is compatible with cases in which the subject is in a position to reasonably believe that she knows p , but not in a position to know that she knows p , due to the fact that that belief is unsafe.

It seems uncontroversial that one can be blameless even if one's second-order belief falls short of knowledge, provided that the belief meets minimal rationality demands. An obvious example is provided by the barn façade case, in which a failure of knowledge is due to a modal environmental condition. Most philosophers agree that in such cases the subject's belief is justified, or at least reasonable and excusable, and therefore it is not worthy of blame. Prudential blameworthiness seems not to be related to the failure of external conditions on knowledge, such as the lack of safety or sensitivity. Rather, for being blameless to ϕ in this sense, many hold that it is sufficient to have reasonable belief that one should ϕ , though one shouldn't ϕ . A subject who acts on what she reasonably believes is fully excusable, and thus not blameable as imprudent.

If this is correct, assuming KNP, the subject cannot be considered imprudent in cases in which she reasonably believes that she knows q but doesn't know that she knows q . Since a failure of the safety condition doesn't undermine any of the conditions for having a reasonable second-order belief, the non-luminosity of knowledge in high stakes has no direct relevance to whether one is blameable as imprudent or not. So a subject in a high-stakes situation who reasonably believes that she knows q is prudentially blameless in holding the second-order belief and in using q as a premise in her practical reasoning, even though her second-order belief is unsafe (and thus the subject doesn't know that she knows q). Therefore, the reason why luminosity fails for knowledge

(according to the anti-luminosity argument) doesn't explain why S^* would be blameworthy if she were relying on p in practical reasoning.¹⁷

In the case of S^* , our intuition that it would be blameworthy for S^* to treat p as a reason in practical reasoning is not due to her lack of second-order knowledge. At most it could be due to the fact that S^* cannot reasonably believe that it is appropriate to rely on p . But Williamson's anti-luminosity argument doesn't show that in cases in which luminosity fails, one is also not in the position to reasonably believe oneself to be in the relevant condition.

It might be argued that in typical high-stakes cases exemplified in the literature, should the subjects believe that their first-order beliefs constitute knowledge, those second-order beliefs would be not only unsafe, but also unreasonable, and thus blameworthy. These cases are set up in such a way that it seems clear to the subjects themselves that the evidential grounds for their first-order beliefs are not very robust, not enough to possess iterated knowledge of the relevant proposition. If these subjects were considering whether they know the relevant proposition, the reasonable attitude for them to take would be agnosticism. Thus, although we can in principle separate the belief's properties of being unreasonable and of being unsafe from error, in practice it seems very hard to distinguish them in those high stakes cases. As a result, in all these cases the subjects would be considered blameworthy for not suspending judgment on whether they know the relevant propositions.¹⁸

Nonetheless, even admitting that lack of reasonability and of safety from error cannot be easily separated in standard high-stakes cases, the point still remains that in those high-stakes cases the blameworthiness for relying on p in practical reasoning is due to the unreasonableness, not the unsafety, of the second-order belief that p . Williamson's account would still be wrong in virtue of

¹⁷ Similar arguments apply to other iterations of knowledge as well. By discriminating reasonable belief from safe belief, we can see that iterations of knowledge are not directly related with prudential blameworthiness.

¹⁸ I am indebted to a referee of this journal for encouraging me to address this point.

appealing to the wrong *explanans* (viz., anti-luminosity and lack of safety), even if it were eventually able to provide the right verdict for most high-stakes cases.

Furthermore, it is not difficult to find high-stakes cases in which it seems fully reasonable for the subject to believe that one knows the relevant proposition, even though that belief is unreliable. Williamson's account fails to provide the correct verdicts about these cases, since the lack of second-order knowledge doesn't prevent the reasonability and blamelessness of holding second-order belief and of relying on the relevant proposition in practical reasoning.

Consider a specific example similar to the case of feeling cold in the original anti-luminosity argument. Suppose that an investor allocates a certain capital in an investment with the expectation of a future financial return. The return on her investment has the property of growing slowly but stably in value when economic circumstances are favourable, but of collapsing below the initial value in unfavourable circumstances. Though circumstances are favourable now, a downturn might occur in the near future. Assume that the investor is a normal risk-averse subject: her marginal utility decreases with an increase in monetary value. The investor is reasonably disposed to keep her investment until it passes a certain threshold, after which it would be too risky not to sell (i.e., the **expected** disutility of the loss would be higher than the **expected** utility of the gain). Suppose also that throughout the process the investor regularly considers whether she knows that the investment is not risky at the moment. The case is designed in such a way that the risk of not selling the investment grows very slowly, almost imperceptibly, but steadily, as happens in Williamson's original case with the feeling of heat. As in that case, the investor's power of discriminating the value is limited: the investor cannot discriminate any significant change in risk within a few hours (though she can over several days). The subject is very confident that the investment is not risky at the beginning, and then her confidence gradually diminishes with the growth of risk. In this respect, the subject's confidence at each stage of this process should be considered as reasonable as that of the subject in Williamson's original case. An application of the

anti-luminosity argument shows that there is a time t at which the subject's belief that she knows that the investment is not risky is true and reasonable but unsafe, i.e. such that it is too close to a case in which the subject loses knowledge that the investment is not risky.¹⁹ In this case, it seems that at time t the subject is not blameworthy for taking herself to know that the investment is not risky and for not selling. The situation is analogous to the original 'feeling cold' case, where the subject is not unreasonable or blameworthy for believing that she feels cold when this belief is unsafe just for an *indiscriminable* margin. Nonetheless, at time t our investor lacks knowledge that she knows that the investment is not risky. Williamson's account predicts the wrong result about this case, since the lack of second-order knowledge doesn't prevent the reasonability and blamelessness of holding second-order belief that one knows and of relying on the relevant proposition in practical reasoning.²⁰ This type of case shows that Williamson's account is doubly wrong: not only does it identify the *explanans* with the wrong property (iterated knowledge instead of reasonable belief), but it also fails to predict rational behaviour in a range of cases in which conditions for blameworthiness come apart from those for safety.

References

- Broome, John. 2013. *Rationality Through Reasoning*. Wiley-Blackwell.
Brown, Jessica. 2008. 'Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning'. *Noûs* 42 (2): 167–89.

¹⁹ Note that in the close case in which the subject loses knowledge that the investment is not risky, the loss of knowledge is due to the unsafety of the belief that the investment is not risky, not the falsity of that belief. Otherwise, at time t , the subject would already lose knowledge that the investment is not risky since her belief in the relevant proposition at time t would be unsafe.

²⁰ Let me just observe here that this example is a high-stakes case in the sense used in the relevant literature (e.g., Fantl and McGrath 2002; Stanley 2005; Weatherson 2005), one in which a lot turns on whether the relevant proposition (that the investment is not risky) is true. The reader should be careful not to confuse the investment's risk (measured by the ratio between the disutility of the loss and the utility of the gain) with the stakes on whether the investment is not risky. While the former grows with time, the latter are high at all times in the process.

- Cohen, Stewart. 1999. 'Contextualism, Skepticism, and the Structure of Reasons'. *Philosophical Perspectives* 13 (s13): 57–89.
- Crisp, Thomas M. 2005. 'Hawthorne on Knowledge and Practical Reasoning'. *Analysis* 65 (286): 138–40.
- DeRose, Keith. 1992. 'Contextualism and Knowledge Attributions'. *Philosophy and Phenomenological Research* 52 (4): 913–29.
- . 2009. *The Case for Contextualism*. New York: Oxford University Press.
- Fantl, Jeremy, and Matthew McGrath. 2002. 'Evidence, Pragmatics, and Justification'. *Philosophical Review* 111 (1): 67–94.
- . 2009. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Gerken, Mikkel. forthcoming. *On Folk Epistemology*. Oxford University Press.
- . 2011. 'Warrant and Action'. *Synthese* 178 (3): 529–47.
- . 2015. 'The Role of Knowledge Ascriptions in Epistemic Assessment'. *European Journal of Philosophy* 23 (1): 141–61.
- Greco, Daniel. 2014. 'Could KK Be OK?' *Journal of Philosophy* 111 (4): 169–97.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, John, and Jason Stanley. 2008. 'Knowledge and Action'. *Journal of Philosophy* 105 (10): 571–90.
- Joyce, James M. 2005. 'How Probabilities Reflect Evidence'. *Philosophical Perspectives* 19: 153–78.
- Kinderman, Peter, Robin Dunbar, and Richard P. Bentall. 1998. 'Theory-of-Mind Deficits and Causal Attributions'. *British Journal of Psychology* 89 (2): 191–204.
- Markovits, Julia. 2011. 'Internal Reasons and the Motivating Intuition'. In *New Waves in Metaethics*, edited by Michael Brady, 141–65. Palgrave Macmillan.
- McKenna, Robin. 2015. 'Assertion, Complexity, and Sincerity'. *Australasian Journal of Philosophy* 93 (4): 1–17.
- Reed, Baron. 2010. 'A Defense of Stable Invariantism'. *Noûs* 44 (2): 224–44.
- Stanley, Jason. 2005. 'Fallibilism and Concessive Knowledge Attributions'. *Analysis*, no. 65: 126–31.
- Stiller, James, and R. I. M. Dunbar. 2007. 'Perspective-Taking and Memory Capacity Predict Social Network Size'. *Social Networks* 29 (1): 93–104.
- Weatherston, Brian. 2005. 'Can We Do without Pragmatic Encroachment?' *Philosophical Perspectives* 19 (1): 417–43.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- . 2005. 'Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge'. *The Philosophical Quarterly* 55 (219): 213–35.