

Margin for error and the transparency of knowledge

Paul Egré, Jérôme Dokic

► **To cite this version:**

Paul Egré, Jérôme Dokic. Margin for error and the transparency of knowledge. v2, 2008, pp.1-36.
ijn_00000479v2

HAL Id: ijn_00000479

https://jeannicod.ccsd.cnrs.fr/ijn_00000479v2

Submitted on 23 Aug 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Margin for Error and the Transparency of Knowledge*

Jérôme Dokic (EHESS, IJN)

dokic@ehess.fr

Paul Égré (CNRS, IJN)

paulegre@magic.fr

.....

Abstract. In chapter 5 of *Knowledge and its Limits*, T. Williamson formulates an argument against the principle (KK) of epistemic transparency, or luminosity of knowledge, namely “that if one knows something, one knows that one knows it”. Williamson’s argument proceeds by *reductio*: from the description of a situation of approximate knowledge, he shows that a contradiction can be derived on the basis of principle (KK) and additional epistemic principles that he claims are better grounded. One of them is a reflective form of the margin for error principle defended by Williamson in his account of knowledge. We argue that Williamson’s *reductio* rests on the inappropriate identification of distinct forms of knowledge. More specifically, an important distinction between perceptual knowledge and non-perceptual knowledge is wanting in his statement and analysis of the puzzle. We present an alternative account of this puzzle, based on a modular conception of knowledge: the (KK) principle and the margin for error principle can coexist, provided their domain of application is referred to the right sort of knowledge.

.....

* Penultimate draft (August 2007). Forthcoming in *Synthese*.

In chapter 5 of *Knowledge and its Limits*, Williamson formulates an argument against the principle (KK) of epistemic transparency, or luminosity of knowledge, namely “that if one knows something, then one knows that one knows it” (Williamson 2000: 115). Principle (KK), which corresponds to axiom schema 4 of propositional modal logic, is also called “positive introspection” and was originally defended by Hintikka in his seminal work on epistemic logic (Hintikka 1962: c. 5, “Knowing that one knows”). Williamson’s argument proceeds by *reductio*: from the description of a situation of approximate knowledge, he shows that a contradiction can be derived on the basis of principle (KK) and additional epistemic principles that he claims are better grounded. One of them is a reflective form of the margin for error principle defended by Williamson in his account of knowledge. We argue that Williamson’s *reductio* rests on the inappropriate identification of distinct forms of knowledge. More specifically, an important distinction between perceptual knowledge and non-perceptual knowledge is wanting in his statement and analysis of the puzzle. The (KK) principle and the margin for error principle can coexist, provided their domain of application is referred to the right sort of knowledge.

1. Margin for error

Williamson’s argument against positive introspection rests on three other epistemic principles: a margin for error principle, knowledge of the margin for error principle, and a principle of closure. As we will see shortly, only the last two principles are actually needed for Williamson’s *reductio*, but since knowledge of the margin for error principle is parasitic on the margin for error principle itself, it is relevant to make the distinction. The margin for error principle is undoubtedly the main originality in Williamson’s theory of knowledge and the principle calls for a preliminary explanation before we state his argument.¹ The margin for

error principle is intended to express a reliability condition for knowledge. As Williamson puts it (2000: 17):

Where one has only a limited capacity to discriminate between cases in which p is true and cases in which p is false, knowledge requires a margin for error: cases in which one is in a position to know p must not be too close to cases in which p is false, otherwise one's beliefs in p in the former cases would lack a sufficiently reliable basis to constitute knowledge.

Formally, the margin for error principle can be seen as a generalization of the principle of factivity of knowledge to the notion of approximate knowledge. To take a specific example: the factivity principle says that to know that someone with i hairs is bald implies that someone with i hairs is indeed bald. Assuming a fixed margin of 1 hair, the margin for error principle says that to know that someone with i hairs is bald implies that someone with $i+1$ hairs and someone with $i-1$ hairs are also bald (for $i > 0$). In other words, it is impossible to know that someone with i hairs is bald, while someone with $i+1$ hairs is not bald, because knowledge is necessarily approximate in such a case. Knowledge, according to Williamson, is not merely factive, we may say that it is factive with respect to neighbouring cases. In this situation, the principle takes the following form (where $bald(i)$ means that someone with i hairs on their head is bald):

$$(a) K(bald(i)) \rightarrow (bald(i-1) \& bald(i) \& bald(i+1))$$

In section 5.1 of *Knowledge and its Limits*, Williamson states a particular form of the principle for the evaluation of the height of a tree. The logical form of the principle is different there. Suppose a certain tree is i inches tall. For every natural number i , let p_i represent the statement that the tree is i inches tall. In this situation, we cannot have:

$$(b) Kp_i \rightarrow (p_{i-1} \& p_i \& p_{i+1})$$

because it would mean: to know that the tree is i inches tall, the tree has to be $i-1$ inches tall and i and $i+1$ inches tall, a plain contradiction. So, the margin for error principle takes a negative form: “for no natural number i is the tree $i+1$ inches tall while [Mr Magoo] knows that it is not i inches tall” (Williamson 2000: 115). Hence, the following schema is assumed for all i :

$$(c) \neg(p_{i+1} \& K\neg p_i)$$

which is equivalent to:

$$(d) K\neg p_i \rightarrow \neg p_{i+1}$$

The situation being symmetric with respect to $i-1$, we get as a counterpart to (a):

$$(e) K\neg p_i \rightarrow (\neg p_{i-1} \& \neg p_i \& \neg p_{i+1})$$

that is, Mr Magoo knows that the tree is not i inches tall only if the tree is not $i-1$ or i or $i+1$ inches tall. More generally, when the margin is n , the general version of (a) becomes: for all k such that $|i-k| \leq n$, $K(bald(i)) \rightarrow bald(k)$, and likewise instead of (e) we will have that for every k such that $|i-k| \leq n$, $K\neg p_i \rightarrow \neg p_k$.

Margin for error principles play a crucial role in Williamson's version of the epistemic theory of vagueness, since they give an explanation of why one may not be able to judge whether someone is bald or not: for instance, suppose the cut-off for baldness is 4000 hairs. If one's margin for error is 1000 hairs, it follows that one will not be able to know that someone with 3500 hairs on their head is bald, or that someone with 4700 hairs on their head is not bald. Likewise, when a character like Mr Magoo makes judgements about the height of a tree, his judgements involve a margin for error: for Magoo to be sure that the tree is less than a certain height, a sufficient margin for error is needed, below which his judgements start to lose their reliability. Although one may not necessarily want to follow Williamson in his epistemic account of vagueness, we think that his use of the margin for error principle is fairly plausible in the kind of scenario that he discusses in order to refute the (KK) principle, namely situations in which one is to give explicit estimates about a certain quantity, based on one's qualitative perception. Since Williamson's rebuttal of (KK) hangs on accepting the margin for error principle, this will explain our attempt to preserve both principles.

2. Williamson's puzzle

Williamson's puzzle may be described as an *epistemic sorites*: Mr Magoo observes a certain tree and makes judgements about its height; from the assumption that Magoo's knowledge is positively introspective, and that Magoo knows that his knowledge involves a margin for error, we end up with the counterintuitive conclusion that Magoo's knowledge can be extended indefinitely. The puzzle can be formalized using a propositional modal language. The usual rules of *modus ponens* and substitution are assumed, and the following axioms and rule are taken to hold for every proposition ϕ and natural number i :

(KME) $K(K\neg p_i \rightarrow \neg p_{i+1})$

(KK) $K\phi \rightarrow KK\phi$

(C) If ϕ follows logically from a set of propositions Γ , such that for all members ψ of Γ , $K\psi$ holds, infer $K\phi$

Define (ME) to be the schematic version of the margin for error principle presented as (d) above (namely the upward part of the schema, for a margin of 1 unit). (KME) states Mr Magoo's knowledge of (ME), namely the fact that "Mr Magoo reflects on the limitations of his eyesight and ability to judge heights" (Williamson 2000: 115). It is therefore a principle of reflection on the margin for error principle. Principle (C) is a principle of closure of knowledge under logical consequence. To avoid too much idealization, Williamson assumes in his presentation that both principles (C) and (KK) are restricted to propositions pertinent to the argument. This restriction will be left implicit here.

The assumption made by Williamson is that Mr Magoo knows that the tree is not i inches tall for a certain number i . We thus have the following derivation:

(1) $K\neg p_i$, hypothesis

(2) $K(K\neg p_i \rightarrow \neg p_{i+1})$, by (KME)

(3) $KK\neg p_i$, by (1) and (KK)

(4) $K\neg p_i, K\neg p_i \rightarrow \neg p_{i+1} \vdash \neg p_{i+1}$, by propositional reasoning

(5) $K\neg p_{i+1}$, by (2), (3), (4) and (C)

The derivation yields as a derived rule that, from $K\neg p_i$, one can infer $K\neg p_{i+1}$. This consequence is puzzling and undesirable because if, for instance, Mr Magoo knows that $K\neg p_0$, then we can infer $K\neg p_n$, for any n , by n applications of the derived rule and *modus ponens*. Thus, supposing Mr Magoo knows that the tree he sees is not 0 inches tall, he knows that it is not k inches tall for a particular k . But if the tree happens to be precisely k inches tall, then Mr Magoo knows something false! In other words, we get a formal contradiction if, besides principles (KME), (KK) and (C), we explicitly assume factivity of knowledge, $K\neg p_0$, and p_k for some k .

How did the puzzle arise? Williamson considers that principle (KME) is well motivated for all i , and that principle (C), seen as a restricted closure principle, is legitimate too, assuming that Magoo makes careful deductions on the basis of what he knows. As a consequence, he takes (KK) to be the culprit. Several other options can be entertained, we should note: to abandon factivity, to reject the premise that the tree has a definite size k , or to deny that Mr Magoo knows the tree is not of size 0. The first two options are clearly off the mark, and the assumption that, for a given tree, the agent knows that it is not of null size seems perfectly harmless: if Mr Magoo can recognize the tree as a tree, then obviously he has to ascribe a positive size to it. Our claim, therefore, is that Williamson goes too quickly over the examination of principles (KK), (KME), (C) and their interaction. We still think that the principle of margin for error makes sense, but there is a specificity of the reflection condition embodied in (KME) that is not taken into account in Williamson's interpretation.

3. Reflective knowledge and perceptual knowledge

In our view, Williamson, in his statement of (KME), fails to take into account the fact that different sorts of knowledge are involved. The same criticism applies to the other principles. We shall argue that the idea that the same sort of knowledge is involved at each level is psychologically implausible. Besides, distinguishing between forms of knowledge can help us to avoid the puzzle without necessarily discarding the motivations underlying the margin for error principle.

In order to see this, let us have a closer look at (KME). By (C), (KME) is equivalent to the schema:

$$K(p_{i+1} \rightarrow \neg K\neg p_i)$$

that is, Mr Magoo knows that if the tree is $i+1$ inches tall, then he does not know that it is not i inches tall. The question we want to raise concerns the reading of the operator K in this sentence. The non-iterative principle (ME), namely $(p_{i+1} \rightarrow \neg K\neg p_i)$, states a constraint concerning Mr Magoo's perceptual knowledge: if the tree is $i+1$ inches tall, then Magoo cannot rule out, simply on the basis of his perception, that it is i inches tall. As Williamson emphasizes, the actual length of the measurement unit does not matter, since one could make the case "even stronger by reducing the interval of an inch to something much smaller" (2000: 115). What (ME) says is that I am not necessarily able to discriminate perceptually between two adjacent sizes, when the interval is sufficiently small. To make explicit the fact that the knowledge involved in (ME) is perceptual, one may therefore write the principle more explicitly as: $(p_{i+1} \rightarrow \neg K_\pi \neg p_i)$, where K_π denotes perceptual knowledge.

In the case of (KME), by contrast, it is doubtful that the outermost knowledge operator should have the same meaning as the embedded one. Let us suppose, for the sake of the argument, that it is the case, and that $K(p_{i+1} \rightarrow \neg K\neg p_i)$ is to be read uniformly as :

$$K_{\pi}(p_{i+1} \rightarrow \neg K_{\pi}\neg p_i)$$

Then it would mean that I know by perception a certain conditional fact about my perceptual knowledge. This is certainly a very controversial claim. Can I really perceive a conditional fact of this kind, in the same sense in which I perceive the height of a tree? Assume, on the other hand, that the knowledge involved in (KME) is some form of non-perceptual knowledge, for which one may introduce a distinct operator K_{ρ} . In that case, the explicit form of the principle becomes:

$$K_{\rho}(p_{i+1} \rightarrow \neg K_{\rho}\neg p_i)$$

In this case, however, the margin for error principle: $p_{i+1} \rightarrow \neg K_{\rho}\neg p_i$, does not necessarily hold. Suppose K_{ρ} is a form of logical knowledge. In the case where $i=0$, principle (ME) would mean: if a quantity is strictly positive, I do not know logically that it is non-zero. Why should it be the case? *A fortiori*, why should (KME) hold, that is why should I know *logically* that if a quantity is strictly positive, I do not know logically that it is non-zero? Logical knowledge is precisely the kind of exact knowledge to which margin for error principles should not apply.

As a consequence, the only plausible construal of the principle (KME) seems to be one in which the embedded operator refers to a form of approximate knowledge. Let us suppose it is perceptual knowledge. The principle then reads:

$$K(K_{\pi} \neg p_i \rightarrow \neg p_{i+1})$$

What is the status of the outermost knowledge operator then? We argued it is not perceptual knowledge, but we think it is a form of *reflective knowledge*, namely non-perceptual knowledge about one's own (in this case, perceptual) knowledge. The crucial fact is that reflective knowledge is not necessarily subject to a margin for error principle. Why should all forms of knowledge be subjected to such a constraint indeed?

More generally, our claim is that the principles (KME) and (KK), which involve iterations of knowledge, should be rephrased in terms of two operators, which we write K_{π} (for perceptual knowledge) and K (for non-perceptual, reflective knowledge). Restricting principle (C) to the operator K , we get as a logic:

$$(KME') K(K_{\pi} \neg p_i \rightarrow \neg p_{i+1})$$

$$(KK') K_{\pi} \phi \rightarrow KK_{\pi} \phi$$

(C') If ϕ follows logically from a set of propositions Γ , such that for all members ψ of Γ , $K\psi$ holds, infer $K\phi$

(KME') now states that it is known non-perceptually that perceptual knowledge obeys the margin for error principle. (KK') now states that if I perceive a certain fact, I know non-perceptually that I perceive it. Principle (C') now specifies that non-perceptual knowledge is closed under logical consequence.

Let us see what happens when we rephrase the puzzle in bimodal terms. We then get:

- (1) $K_{\pi} \neg p_i$, hypothesis
- (2) $K(K_{\pi} \neg p_i \rightarrow \neg p_{i+1})$, by (KME')
- (3) $KK_{\pi} \neg p_i$, by (1) and (KK')
- (4) $K_{\pi} \neg p_i, K_{\pi} \neg p_i \rightarrow \neg p_{i+1} \vdash \neg p_{i+1}$, by propositional reasoning
- (5) $K\neg p_{i+1}$, by (2), (3), (4) and (C')

The principles now yield as a derived rule that if I perceive that the tree is not i inches tall, then I know, non-perceptually, that it is not $i+1$ inches tall. This now seems a safe rule, for if we assume that $K_{\pi} \neg p_0$ holds, we can only infer $K\neg p_1$, without propagating this knowledge to further values. In addition, it is safe in the sense that it is no more puzzling than principle (ME) for perceptual knowledge (remember that the margin for error principle, in this case, asserts that $K_{\pi} \neg p_i \rightarrow \neg p_{i+1}$).

The bimodal epistemic logic we have appealed to here is minimal in that we have only needed to assume thus far principle (KK') (which is a weak version of (KK)), principle (C) for the operator K , the margin for error principle for K_{π} (ME), and that this principle is K-known to be so (principle (KME')). Principles like (KME') and (KK') allow us to spell out plausible interactions between the different forms of knowledge at stake. But more can

reasonably be assumed about the separate and joint behaviour of each of K and K_π . For instance, it is reasonable to suppose that both are factive (which is axiom schema **T** in modal logic), and that their factivity is K -known. Let us consider some other appropriate principles.

Note, first, that we need not dispense with principle (KK), if we think this principle now holds of reflective knowledge only. The reason why we described the knowledge expressed by K as *reflective* knowledge is that it expresses meta-knowledge of some kind, here knowledge about one's perceptual knowledge, as axiomatized by (KME') and (KK'). It is natural, however, to suppose that meta-knowledge about one's reflective knowledge is still reflective knowledge. Let us define a knowledge operator as *positively introspective* if it obeys positive introspection. The adoption of (KK) would mean that the reflective knowledge involved in (KME') and (KK') is also a form of positively introspective knowledge. Again, this seems a natural assumption to make about K , even though there might be forms of reflective knowledge that are not positively introspective.

If positive introspection is adopted for K , one may actually strengthen principle (C') into the distribution axiom **K** in order to recover a full **S4** logic for K . What logic the operator K_π might have, on the other hand, is a more delicate issue that would call for a separate treatment. The central point for our argument, however, is the existence of clear cases where meta-knowledge about one's perceptual knowledge cannot be of the perceptual kind. An iterative axiom like (KME), we argued, cannot meaningfully hold of K_π .ⁱⁱ The argument can be rephrased in a simpler way if, instead of the margin for error principle, we consider the factivity principle. Remember that factivity can be seen as a particular case of the margin for error principle. Likewise, the knowledge that one's knowledge is factive is analogous to the

knowledge that one's knowledge is subject to the margin for error principle. Consider an axiom like (KF), which expresses knowledge that knowledge is factive:

$$(KF) K(K\phi \rightarrow \phi)$$

An iterative axiom of this kind cannot hold of the operator K_π , for $K_\pi(K_\pi\phi \rightarrow \phi)$ would state that I know perceptually that my perceptual knowledge is factive. Yet how could this piece of knowledge be of a perceptual kind? More generally, one may doubt the existence of pure iterative axioms for K_π .ⁱⁱⁱ Just as it is natural to assume that reflective knowledge is also positively introspective because it is essentially a form of iterative knowledge, it is natural to suppose that perceptual knowledge will not be positively introspective if it is not iterative in the first place. If we reject positive introspection for K_π in favour of an axiom like (KK'), it is therefore on conceptual grounds that are very distinct from Williamson's criticism of the luminosity of knowledge.

One may wonder here whether the introduction of distinct knowledge operators, to account for distinct forms or methods of knowledge, does not put into question the notion of the unity of knowledge. We think it need not be so, so long as we can account for the way these methods interact with each other. This is essentially what axioms like (KME') and (KK') perform in our analysis of Williamson's puzzle. But further interaction axioms can be added. In particular, just as (KME') is a bimodal analogue of (KME), there should be a bimodal analogue (KF') of (KF), to express the reflective knowledge that one's perceptual knowledge is factive:

$$(KF') K(K_\pi\phi \rightarrow \phi)$$

Assuming (KK') and (C'), (KF') allows one to infer $K\phi$ from the assumption $K_{\pi}\phi$. The derivation corresponds exactly to the above proof that $K\neg p_{i+1}$ follows from $K_{\pi}\neg p_i$, assuming (KK'), (C') and (KME'). Again this is no surprise, given the analogy between factivity and the margin for error principle. If Kripke's axiom **K** is assumed for K instead of (C'), one can easily prove the stronger result that $K_{\pi}\phi \rightarrow K\phi$, using (KK') and (KF'). In either case, this means that perceptual knowledge can be turned into reflective knowledge. There should be no misunderstanding here about the fact that we characterized reflective knowledge as a form of non-perceptual knowledge. $K_{\pi}\phi \rightarrow K\phi$ should be understood as meaning that knowledge acquired by perceptual means enters as an input to knowledge of a different kind. The only axiom and rule that we want to avoid in our system are the converse $K\phi \rightarrow K_{\pi}\phi$ and $K\phi \therefore K_{\pi}\phi$.^{iv}

At this point, it may be useful to give a more precise characterization of the criteria we consider relevant in order to distinguish between perceptual and reflective knowledge. The first thing to say is that we conceive of both kinds of knowledge as propositional: a subject like Magoo can know visually, for instance, *that the tree in front of him is taller than 2 inches*. To say that he knows this fact *perceptually*, however, is to say that this piece of knowledge is given more or less *directly* through his visual input, without involving routines other than those that are usually at stake in the ordinary perception of distances and relations between objects. For instance, in a given situation, Magoo may be able to *see* that a given object a is taller than another object c , in which case we will say that he knows perceptually that a is taller than c . However, suppose a situation in which Magoo can see that a is taller than b through a certain window, and that he can see through a different window that b is taller than c . Let us suppose, moreover, that Magoo cannot see a and c together. From his memory and

reasoning capacities, Magoo can nevertheless infer and thereby come to know that a is taller than c . In that case, we would not describe the proposition in question as a proposition that Magoo knows *perceptually*, even though he knows the proposition on the basis of his visual perception. The reason is that Magoo's knowledge is not acquired *directly* through vision, but rather *inferred* indirectly from what he sees.^v In our framework, the situation can be represented by assuming the following two premises to hold, namely $K_{\pi}(a > b)$ and $K_{\pi}(b > c)$. Just from those two premises, however, it does not follow that $K_{\pi}(a > c)$, since we made no assumption regarding the closure of K_{π} under logical consequence;^{vi} however, from the derived rule $K_{\pi}\phi \rightarrow K\phi$, it does follow that $K(a > b)$ and $K(b > c)$, and from (C') (which we may suppose to apply to logical consequences based on logical postulates for the relation ">"), it follows in turn that $K(a > c)$. What the example suggests, therefore, is that we do not need to introduce a mixed category between perceptual knowledge and reflective knowledge to account for knowledge based on reasoning from perceptual knowledge. Such a mixed form of knowledge is undeniably needed, but it is already taken care of by the operator K , assuming axioms (KME'), (KK'), (C') and (KF') as background theory.

A second point to note is that the variety of knowledge we call reflective knowledge may very well stand to any more specific kind of knowledge as it stands to what we call perceptual knowledge in the case under discussion. Suppose for instance that Magoo comes to know by testimony that the Pope just died. To circumvent any objection, we may suppose that this information is true, that Magoo will not accept testimonies that are not perfectly reliable, and even that Magoo so far has been infallible in his acceptance and rejection of testimonies. So Magoo knows by testimony that the Pope died. But suppose moreover that Magoo is aware that he knows this information by testimony. Certainly, however, his knowledge that he knows by testimony that the Pope died does not itself count as knowledge by testimony: in normal circumstances, this kind of second-order knowledge is given to him directly by

reflection.^{vii} More generally, therefore, for any specific kind of knowledge, be it knowledge by memory, testimony, a priori reasoning, reflection, and so on, if we represent by $K_?$ the kind of knowledge in question (where the index “?” is dummy for any proper index specifying the relevant kind of knowledge), what we claim is that the operator K of reflective knowledge can stand to $K_?$ as it does to K_π . In particular, we may suppose that the axioms (KK’) and (KF’) hold not only for K *vis a vis* K_π , but also *vis a vis* any specific operator $K_?$. As we saw above, it is reasonable to suppose that one can have reflective knowledge about one’s reflective knowledge, which is why K can be taken to be positively introspective. But if this is so, and if for any indexed operator $K_?$ it therefore holds in the same way that $K_?p \rightarrow Kp$, this means that anything that is known through a specific method can also be known reflectively. Thus, what we call reflective knowledge can also be conceived as a generic kind of knowledge, allowing for sources in all of perception, testimony, and so on. Such a generic kind is needed in any case, in order to maintain an integrated account of knowledge and its different sources.

To summarize what we have said so far: the notion of reflective knowledge encoded by the operator K is more general than that encoded by the operator K_π in three respects. First, unlike K_π , the operator K can embed any other knowledge operator, including itself. Secondly, any proposition that is known by perception can be known by reflection, although the converse does not hold. Finally, this entailment may hold more generally for any specific kind of knowledge, suggesting that reflective knowledge can be conceived as that kind of knowledge that is entailed by any other kind of knowledge. Precisely because of that, however, there is no reason to suppose that the restrictions that apply to specific kinds of knowledge will apply to the reflective-generic kind represented by K . We shall return to this point in section 6, and postpone until there the objections that might be raised against the distinction introduced in this section. Before that, we propose to examine how the thesis of modularity which we defend here bears on the validity of margin for error principles.

4. The Glimpse

In chapter 6 of *Knowledge and its Limits*, Williamson presents a “perceptual” counterpart to the Surprise Examination Paradox, “the Glimpse”, which is structurally analogous to the puzzle stated in section 2 above. He concludes, for the same reasons, that principle (KK) is the problematic assumption at the heart of the paradox. The Glimpse is presented by Williamson in the following manner (2000: 135):

A teacher’s pupils know that she rings all and only examination dates on the calendar in her office. At the beginning of the term, the only knowledge they have of examination dates this term comes from a distant glimpse of the calendar, enough to see that only one date is ringed and that it is not very near the end of the term, but not enough to narrow it down much more than that. The pupils recognize their situation. They know now that for all numbers i , if the examination is $i+1$ days from the end of term then they do not know that it will not be i days from the end ($0 \leq i < n$). In particular, they know now that if it is on the penultimate day then they do not know now that it will be on the last day. But they also know from their glimpse of the calendar that it will not be on the last day. They deduce that it will not be on the penultimate day. They also know that if it is on the antepenultimate day then they do not know that it is not on the antepenultimate day. And so on. They rule out every day of term as a possible date for the examination.

The structural analogy with the puzzle stated in section 2 should be clear. Let p_i now stand for: “the examination is i days from the end of the term”, or equivalently here: “the date ringed on the calendar is i positions from the end-of-term date on the calendar”. Let $K\phi$ represent the statement that the pupils know that ϕ (by whatever method). The assumptions made by Williamson are:

$$K\neg p_0$$

$$K(p_{i+1} \rightarrow \neg K\neg p_i), \text{ for all } i < n \quad (\text{KME})$$

Exactly as in section 2, this entails $K\neg p_1$ under the assumptions (KK) and (C), and by repeated application of this inference pattern, this yields $K\neg p_n$ for every relevant number n .

The argument of Williamson against the principle (KK) can be challenged exactly as previously, by distinguishing between different methods of knowledge. We can still accept the validity of a margin for error principle (ME) in the case of knowledge acquired by glimpsing, as Williamson does. But a bimodal formulation of its epistemic version (KME), in the form of a principle like (KME'), seems more adequate, because knowledge by glimpsing is not the only form of knowledge at stake here. Supplementing (KME') with principles (KK') and (C') as above, we will infer from the fact that the pupils know from their visual experience that the exam does not occur on the last day, that they know reflectively that it does not occur on the penultimate day either, *even though they are unable to ascertain visually that there is no ring on the corresponding date*. Altogether, the pupils know from their glimpse that the exam will not take place on the last day, and they know from a certain reflection on their visual abilities that it cannot take place on the penultimate day either. But it does not rule out the possibility that the exam takes place on the antepultimate day, indeed they cannot know this just from their glimpsing. Note that there is a difference with the Surprise Examination Paradox here. In the Surprise Paradox, the examination is announced to be a *surprise*, and it is this fact which sustains the backwards induction to the conclusion that there cannot be such an examination. In the Glimpse, the pupils know (from their glimpse)

that there will be an exam. But it is not announced to be a surprise; they are simply unable to locate its occurrence precisely on the calendar.

5. Non-perceptual knowledge without a margin for error

Williamson's argument against (KK) trades on a margin for error principle which he applies univocally to at least two different forms of knowledge, namely perceptual knowledge and non-perceptual, reflective knowledge. In contrast, we think that margin for error principles are relative to *specific* discriminative capacities, such as perception. It is implausible to claim that knowledge always involves discriminative capacities *of the same kind*.^{viii} There might be forms of non-perceptual knowledge which do not bring about the kinds of margin for error brought about by perceptual knowledge.

As an example of reflective transitions which clearly do not introduce additional margins for error, let us consider so-called "ascent routines". In Gordon (1995)'s terminology, an ascent routine is a perfectly reliable method of self-ascription of beliefs. In some cases, one can know that one believes that it is going to rain by asking oneself whether it is going to rain.^{ix} This method does not rest on a form of inner perception of one's beliefs. Typically, it is world-directed rather than self-directed. In this sense, it is not an introspective method, even though it yields correct self-ascriptions of belief.

Ascent routines can also ground self-ascriptions of visual experience. For instance, the move from one's visual experience that it is going to rain to the judgment "I seem to see that it is going to rain" is perfectly reliable. Whenever one has the visual experience that *p*, one is *ipso facto* in a position to judge that one has such an experience. Moreover, it is plausible that such a routine enables the subject to gain *knowledge* of the fact that she has the visual

experience that p . Ascent routines are, at least in some cases, methods of non-perceptual self-knowledge.

Now, the use of an ascent routine need not bring about a margin for error. In the case of vision, the margin for error principle requests that cases sufficiently similar to cases in which I visually know that p be cases in which p is true.^x Visual knowledge is a discriminative capacity which plausibly introduces a specific margin for error. However, self-knowledge based on an ascent routine need not introduce an *additional* margin for error. There does not seem to be a requirement that cases sufficiently similar to cases in which I know that I have the visual experience that p are cases in which I have the visual experience that p . The relevant ascent routine (from one's experience that p to the corresponding judgment) is perfectly reliable precisely because it is not based on the exercise of a fallible capacity such as introspection as traditionally conceived.

An adequate discussion of the epistemology of ascent routines is beyond the scope of this paper. At least two issues should be discussed further. First, in the foregoing example, the ascent routine works only if the subject is able to separate what is visually given from any extraneous (conceptual) information (see again Evans 1982). In the case of the Müller-Lyer illusion, one can know that one has the visual experience as of two unequal lines even if one also knows that the lines are in fact equal. The latter piece of knowledge is extraneous to the content of one's visual experience, which is relatively modular. (We'll return to the modularity of perception in the next section.) Second, ascent routines are available only from the content of experience as it is apprehended by the perceiver. Thus, the ascent routine from seeing (what is in fact) an armadillo to judging "I seem to see an armadillo" is perfectly reliable, but cannot be used by the perceiver if she does not know what an armadillo is, or

does not know what one looks like. By contrast, if she sees something *as* an armadillo, she can move directly from her first-order experience to the higher-order judgment that she has a visual experience as of an armadillo without bringing about an additional margin for error.^{xi}

It is important to note, though, that our argument does not hang on the existence of forms of knowledge which do not bring about a margin for error. It is enough for us to show that perceptual knowledge and non-perceptual knowledge do not introduce the same kind of margins for error. Once again, if Mr Magoo knows that the tree is not i inches tall by visual means, then he can reason to the conclusion that it is not $i+1$ inches tall. This reflective piece of knowledge may or may not bring about a margin for error additional to the margin for error brought about by Mr Magoo's perceptual knowledge. Suppose for the sake of argument that it does. Then cases sufficiently similar to cases in which Mr Magoo knows by reasoning that the tree is not $i+1$ inches tall are cases in which the tree is not $i+1$ inches tall. However, the standard of similarity involved here is specific to the reflective method of knowledge used by Mr Magoo. In particular, the case in which the tree is not $i+2$ inches tall may no longer be relevantly similar to the cases in which Mr Magoo uses this method. Williamson's puzzle still does not arise: Mr Magoo may not be able to reason to the conclusion that the tree is not $i+3$ inches tall from his initial perceptual knowledge that it is not i inches tall.

To our minds, margin for error principles are best conceived as relative to *methods* or *forms* of knowledge, namely those that involve some *specific* discriminative capacity. From the premise that Mr Magoo visually knows that the tree is not i inches tall, one may indeed conclude that he knows through reasoning that the tree is not $i+1$ inches tall (which he may also know visually, of course). Now suppose that there is a precise point n at which Mr Magoo does not visually know that the tree is not n inches tall but does visually know that it

is not $n-1$ inches tall. On our account, it entails that Mr Magoo can still know through reasoning that the tree is not n inches tall. There is no paradox here: the fact that a truth is not known visually is compatible with the fact that it is known by other, non-perceptual means. The margin for error brought about by visual perception need not concern knowledge gained by reasoning, and thus one can solve Williamson's puzzle without necessarily impairing the validity of the (KK) principle.^{xii}

6. Reply to some objections

In the previous sections we have argued that Williamson's attack against the principle of positive introspection rests on an equivocation between two types of knowledge, namely a reflective type of knowledge, susceptible of iterations, but not necessarily involving a margin for error, and a perceptual type of knowledge, involving a margin for error, but most likely not susceptible of iterations. In this final section of the paper, we turn to the discussion of several objections that may be raised against the present account.

1. First, note that what we are claiming is not that Williamson fails to acknowledge that there are different sources and types of knowledge in general. The idea that knowledge comes in different varieties is fairly uncontroversial, and this is certainly not a view that Williamson would deny. Rather, the core of our argument is that Williamson fails to take into account a dimension of modularity of knowledge in the scenario he is describing: the knowledge I have about my visual knowledge need not be constrained in the way in which my visual knowledge is constrained. Nevertheless, Williamson may still resist our criticism, by pointing out that the limitations inherent to the subject's perceptual knowledge are inherited by higher-order levels of knowledge. Thus, in his statement of the puzzle, Williamson considers that "Mr Magoo's

deductive capacities do not fully enable him to overcome the limitations of his eyesight and ability to judge heights” (2000: 116). In particular, the bimodal reformulation of the principles we suggested above should preclude a situation in which Magoo is unable to ascertain visually that the tree is not i inches and yet is able to know deductively (from his visual information) that the tree is not i inches. In other words, Williamson may insist that the conjunction $K\neg p_i \ \& \ \neg K_{\pi}\neg p_i$ is inconsistent – namely, that if it is known reflectively that the tree is not i inches tall, it is also known visually. This would block our solution to the puzzle. But why should this entailment hold in general, other than by a *petitio principii*? Consider again the case of the Müller-Lyer illusion, in which two lines are seen as unequal, while they are in fact equal. This is a situation in which our reflective knowledge that the two lines are of equal size is not backed up by our initial visual perception. The case of visual illusions is probably the best kind of example we can think of in favour of the modularity we claim is relevant in Williamson’s puzzle. More generally, we certainly agree with Williamson that Magoo’s reflective capacities do not “fully enable him to overcome the limitations of his eyesight and ability to judge heights”, but simply from this it does not follow that Magoo’s reflective capacities should be subject to *exactly the same* limitations that affect his visual abilities and the spontaneous judgments that depend on those capacities. A further argument is in any case required to support the idea that higher-order knowledge and first-order knowledge obey the same margins of error.

2. A more specific criticism that may nevertheless be raised against our proposal concerns the prediction made in our system that if Magoo knows visually that the tree is not n inches tall, then he can know reflectively that it is not $n+1$ inches tall. The objection comes in two stages: to begin with, one may accept our claim of modularity, and nevertheless reject the idea that by combining reflection and perception, one can come to know more than by perception

alone. In particular, if Magoo cannot tell by perception that the tree is not $n+1$ inches tall, how can he tell it by reflection plus perception? Wouldn't Magoo need some further empirical input? Our answer to this question is that, from a logical point of view, there is no reason to suppose that if a piece of knowledge is not derivable on the basis of perception alone, it will not be derivable on the basis of perception plus reflection. To claim otherwise is to assume that the addition of reflection to perception is systematically conservative over perception. But it is not necessarily so: as we saw earlier, I can see that an object a is taller than an object b , see that b is in turn taller than a third object c , and infer upon reflection that a is taller than c , without being able to ascertain the fact visually. The situation is similar in the case of the Müller-Lyer illusion which we mentioned above.

Still, and this is the second part of the objection, our system may yield too strong predictions in some cases. An example to that effect was given by J. Snyder (p.c.). Snyder imagined a situation in which Magoo knows perceptually that the tree is 1009, 1010, or 1011 inches tall, something we can symbolize as $K_{\pi}(p_{1009} \vee p_{1010} \vee p_{1011})$, but cannot tell by perception which of those three heights it is. However, he supposes that Magoo can see that the tree is not 1008 inches tall, and also that it is not 1012 inches tall, namely $K_{\pi}\neg p_{1008}$ and $K_{\pi}\neg p_{1012}$. Assuming the version of the rule obtained in our system that would result from a downward version of the margin of error principle, namely $K_{\pi}\neg p_{i+1} \therefore K\neg p_i$, it follows from the last two assumptions that $K\neg p_{1009}$ and $K\neg p_{1011}$, and by (KF') and (C'), it follows from $K_{\pi}(p_{1009} \vee p_{1010} \vee p_{1011})$ that Kp_{1010} : I therefore come to know reflectively that the tree is 1010 inches tall, although by assumption I do not know this perceptually. Here Snyder asks: "Where did the extra information come from? If your perceptual system was not good enough to distinguish between p_{1009} , p_{1010} , and p_{1011} , and if the information you had about the tree came from your perceptual system, then how did you end up knowing p_{1010} , perceptually or not? The problem seems to be your analog of Williamson's derived rule".

We agree with Snyder that the prediction may be too strong in this particular case, granting the plausibility of the premises. To begin with, we have granted Mr Magoo a piece of disjunctive knowledge of the form $K_{\pi}(p_{1009} \vee p_{1010} \vee p_{1011})$, but more realistically, he will be able to see that the tree's height is anywhere on a continuum from 1009 to 1011 inches. So even if Mr Magoo can exclude that the tree is either 1009 or 1011 inches tall, he will be left with a host of remaining options and his knowledge of the tree will not be exact.

Another point is that our derived rule depends on principle (KME') in the first place, namely the analog of Williamson's (KME). It is important to bear in mind that we assumed (KME') for the sake of the argument, namely to show that, even if such a rule is assumed, no pernicious sorites need follow. That being said, we think the real question behind the objection concerns in fact the *origin* of such a principle, namely how the subject comes to know the margin of error of his perceptual knowledge (an assumption we have no reason to deny to rational subjects). In our discussion of Williamson's scenario, we assumed that the subject could reflect on his perceptual abilities, but following Williamson, we did not give an account of the origin of this knowledge. Several answers to the question "Where did the extra information come from?" are conceivable, however. We could imagine that the subject's reflective knowledge of his perceptual limitations is acquired empirically, or even given by some oracle. If so, it is conceivable that the subject comes to know some information that he cannot ascertain simply on the basis of his perception. The prediction, once again, is not by itself implausible: let us think again of the reasoning made by the pupils in the case of the Glimpse. If the pupils are certain that the date ringed was not the last one, and if they can infer, on the basis of their past experience or from what they take to be a regularity of their perceptual apparatus, that this implies the ring could not be on the penultimate day either, then they come to know this fact reflectively, even as the visual stimulus is no longer there to confirm this piece of reflective knowledge. Thus, although the variety of knowledge we call

reflective is tied to some notion of immediate or spontaneous *awareness* through an axiom like (KK') in particular, we consider possible that an axiom like (KME') expresses a form of self-knowledge acquired empirically, more than by inner sense.^{xiii} If such is the case, this will undeniably give support to the idea that the kind of reflective knowledge in question is also inexact and subject to a margin for error. But the point remains that this margin need not be of the same kind as the margin affecting perceptual knowledge.

3. A third and more radical objection is the following. We said in section 3 that the variety of knowledge we call reflective is also generic, in the sense that it can embed any other kind of knowledge, and be such that it is entailed by any other kind of knowledge. Moreover, we have admitted that such a general kind of knowledge is needed, allowing for sources in all of perception, a priori reasoning and so on. But then, as was pointed out to us, maybe Williamson would argue that the principles he is discussing apply to this general kind of knowledge. More explicitly, the objection may be put as follows: if it holds that any proposition that is known perceptually is also known reflectively, and if reflective knowledge is furthermore entailed by any specific kind of knowledge, then why not rephrase the puzzle directly with respect to that general kind of knowledge? Our answer is that this move would not be licit, however, on pain of committing a logical fallacy. To take an analogy first, suppose someone were to draw a general conclusion about mammals from some feature that is specific to dogs only: we would object by saying that “mammal” has been unduly substituted for “dog” somewhere, even though we agree that every dog is a mammal. The situation is analogous here: from the fact that K_{π} entails K (namely from the schema $K_{\pi}\phi \rightarrow K\phi$ which we accept), it does not follow that one can *substitute* the operator K to K_{π} everywhere, for this would mean that any property that holds only of perceptual knowledge also holds of the notion encoded by K , something which is not the case. Our point is precisely

that there are no reasons to think that margin for error principles apply to any kind of knowledge in the same way that they apply to perceptual knowledge. Likewise, we argued that perceptual knowledge can fail to be iterative, but we would not want to ascribe this failure to other kinds of knowledge, and therefore not to the kind of reflective knowledge described by the operator K . In footnote iv, we gave a model-theoretic illustration of this point: while the validity of $K_{\pi}\phi \rightarrow K\phi$ in Kripke frames means that any proposition that is known perceptually is known reflectively, it also implies, conversely, that the accessibility relation underlying K is *included* in the one underlying K_{π} , and therefore that the operator K_{π} can very well fail to be positively introspective, without K failing to be so. Consequently, to insist in this way that Williamson's argument be phrased using only one operator would be to beg the question.

4. Despite this, how should we respond to someone who would insist that the notion of knowledge involved in Williamson's argument is the generic notion of knowledge? Isn't our criticism of Williamson's argument, in that case, an implicit denial of the validity of the *unrestricted* margin for error principle? Our answer is that this unrestricted principle would indeed be implausibly strong, even in the particular scenarios discussed by Williamson. To make the point vivid, let us go back to the example of the Glimpse, and see more carefully how the pupils might reasonably elaborate on their knowledge that the date ringed on the calendar is not the last one. To make things concrete, consider a particular pupil – let us call her Marge – who glimpsed the calendar. All Marge can remember and therefore know initially is that the date ringed was not on the last day of the calendar. Suppose there are 90 days on the calendar. What Marge knows is that the date ringed is not on day 90. In order to represent the notion of knowledge as generic, we will symbolize this as $K\neg p_{90}$, no longer subscripting the knowledge operator in a way that represents specific types of knowledge.

Now, Marge wonders if the date ringed could have been on day 89. But Marge knows that if that were the case, she would not have been able to see that the ring was not on day 90, because the two dates are too close on the calendar for her to discriminate. Thus, Marge can safely infer that the day ringed is not day 89, that is $K\neg p_{89}$. Now, suppose moreover, that the date actually ringed on the calendar is day 88, that is the exam is to take place on the antepenultimate day of the term. Can Marge here infer in the same way that the date ringed cannot be on day 88? We see no reason to suppose so; more than that, this situation is likely to provide a counterexample to the unrestricted margin of error principle. Indeed, if Marge had seen that the date ringed was not 89, in the same way in which she remembers that the date ringed is not 90, then she would be able to infer that the date is not 88. However, by hypothesis, the fact that the date ringed on the calendar is not on 89 is not her initial input. Moreover, day 88 is in principle distant enough from day 90 for Marge to discriminate between them. If so, and if Marge is able to infer that the date ringed is not on day 89, while it is actually on day 88, then this suggests that the margin for error principle $K\neg p_i \rightarrow \neg p_{i-1}$ does not hold in full generality. On our view, this example makes plausible that there is a particular number $i-1$ such that the date is in fact exactly $i-1$, but that Marge knows that the date is not i .

The same can be said about the scenario involving Mr Magoo. The difference with the Glimpse is that, unlike Marge, Mr Magoo has a sustained visual experience of the tree, and can thereby acquire several pieces of knowledge of the form “The tree is not n inches tall” (i.e. $K\neg p_n$). This difference is not substantial, however, for once again, there will be a point at which Mr Magoo can see and safely estimate that the tree is not i inches tall, but cannot see in the same way that the tree is not $i+1$ inches tall, perhaps because the tree is in fact $i+2$ inches tall. Let us suppose that $i = 1000$, and that the tree is 1002 inches tall. On the basis of his visual experience of the tree, Mr Magoo knows that $\neg p_{1000}$. At this point, we can assume that there is a valid application of the unrestricted margin for error principle, namely:

$$K\neg p_{1000} \rightarrow \neg p_{1001}$$

This application is valid because the possibility that p_{1001} is relevant to whether Mr Magoo possesses knowledge that $\neg p_{1000}$. For suppose that p_{1001} were the case. Then Mr Magoo might easily have been mistaken in forming the perceptual judgment that $\neg p_{1000}$, given that he cannot discriminate between a situation in which p_{1000} and a situation in which p_{1001} .

Mr Magoo, being a reflective thinker, knows that $\neg p_{1001}$. Now, is this further application of the unrestricted margin for principle also valid?

$$K\neg p_{1001} \rightarrow \neg p_{1002}$$

It should be obvious that it is not so. For the possibility that p_{1002} is now irrelevant to whether Mr Magoo possesses knowledge that $\neg p_{1001}$ (it would be relevant if Magoo had been in a position to *see* at the beginning that the tree was not p_{1001} , but remember that it cannot be the case). Even if p_{1002} is the case, Mr Magoo remains in a position to know that $\neg p_{1001}$ on the basis of his inference from his visual knowledge that $\neg p_{1000}$. The fact that p_{1002} is no threat to his actual knowledge that $\neg p_{1001}$, in contrast to the previous situation, where the possibility that p_{1001} *was* a threat to Mr Magoo's actual knowledge that $\neg p_{1000}$.

By leaving underspecified the notion of knowledge in these two scenarios, we therefore conclude that the margin for error principle does not hold in full generality. It does not follow, however, that we should entirely forsake it: the point of the previous sections was that if this principle is properly restricted (by distinguishing the methods of knowledge relevant to its applications), then one can do better justice to its initial plausibility.

5. The fallacy we pointed out in 3 above invites us to dispel another potential misunderstanding. The misunderstanding concerns the scope of the (KK) principle. One may be tempted to think that if reflective knowledge is generic in the sense of being entailed by any specific kind of knowledge, and if it obeys positive introspection, then any specific kind of knowledge should satisfy the (KK) principle too, including perceptual knowledge. We explained, however, in what sense this reasoning is flawed. The reason why it is tempting to make this inference, however, is that if *any kind of knowledge* is positively introspective, then it should follow by universal instantiation that perceptual knowledge too is positively introspective. Since in our system K obeys positive introspection, while K_π does not, this implies that the operator K cannot be taken to mean “any kind of knowledge” in this rather loose sense. In this respect, it is important to realize that we agree with Williamson on the following general consequence of his argument, namely the fact that the (KK) principle does not hold unrestrictedly. In Williamson’s account, however, what this means is that knowledge fails to be introspective *tout court*. In our account, by contrast, what this means is rather that some varieties of knowledge may simply fail to be iterative. The point, however, is that the kind of knowledge encoded by the K operator in our system, namely reflective knowledge, can very well be positively introspective and be such that it is entailed by any other kind of knowledge. This makes an important conceptual difference, since the lesson Williamson draws from the incompatibility of positive introspection with margin for error principles is that “we have no cognitive home” (2000, chap. 4). While our aim is not to maintain the universal validity of positive introspection, we may still consistently hold that reflective knowledge remains one of those cognitive homes.

6. Another objection to our argument concerns the idea that perceptual knowledge is not susceptible of iterations, and more generally the criteria that we use to distinguish between so-

called perceptual and non-perceptual knowledge. In the treatment presented above, we insisted indeed that the kind of meta-knowledge one has about one's visual knowledge is not itself visual. But this kind of limitation may seem too strong. Josh Snyder (p.c.) came up with the following counterexample: "At the eye-doctor, examining the eye-charts, I discover that I cannot read the letters below a certain size. This certainty seems to be perceptual knowledge of the limitations of my perceptual capabilities". To make the example concrete, consider a situation in which I see a letter of which I cannot tell whether it is a D or an O (although I am fairly confident it has to be one of them). The situation therefore seems to support the following margin for error principle: $D \rightarrow \neg K_{\pi}\neg O$, namely if it is a D, I do not exclude visually that it is an O, but also the iterative version: $K_{\pi}(D \rightarrow \neg K_{\pi}\neg O)$, namely I know visually that if it is a D, then I might be seeing an O. Our answer, however, is that one ought to distinguish more carefully between the content of one's visual experience and the content of one's judgement. I first see a letter of which I am uncertain whether it is a D or an O. I then make the judgement that if it is a D, I might be seeing an O. Although that judgement about my visual capabilities is acquired on the basis of my visual experience, it is not an item of visual experience properly speaking. For that same reason, we think the logical form of the reflective version of the margin for error principle should be: $K(D \rightarrow \neg K_{\pi}\neg O)$. But once again, what we call reflective knowledge should be thought of as a generic kind of knowledge, not necessarily subject to the same constraints which bear upon more specific forms of knowledge which serve as input to it.

7. The last objection we want to consider concerns theoretical economy. Although one may feel inclined toward the kind of modular treatment of knowledge that we are advocating here, one may find the introduction of two distinct knowledge operators too costly. This move may seem somewhat *ad hoc*, and it threatens the epistemic vocabulary with endless multiplication.

If two operators are allowed, why not carry on with three, and so on, in order to make even finer distinctions (for example between varieties of perceptual knowledge)? Wouldn't it be more elegant to maintain a single knowledge operator, and incorporate modularity at the semantic level rather than the syntactic level? A modular semantics of this kind is conceivable, however, and was actually developed in a sequel to this paper by Bonnay & Egré (2006): instead of using two distinct knowledge operators and putting syntactic restrictions on their iterative behavior and interaction, as we did here, the authors state satisfaction clauses that lead to a differentiated treatment of iterated and non-iterated modalities for an epistemic logic with only one epistemic operator. The conceptual motivation remains essentially the same, however. Here we focussed on the original formulation of Williamson's puzzle, and showed in what way a minimal modification of Williamson's own syntactic assumptions in terms of two knowledge modalities is both more plausible and prevents the paradox. But one may equally well start with a semantic version of Williamson's paradox, adopting the kind of margin for error semantics presented in the appendix to Williamson (1994), and give a "modular" version that validates the principle of positive introspection. Either way, the lesson of Williamson's puzzle is not necessarily that we should get rid of the introspection principles, but rather that more modularity is needed in the way we conceptualize knowledge.

Acknowledgments

The ideas of this paper grew out from the seminar organised by Pascal Engel on Williamson's book *Knowledge and its Limits* in 2003-2004 at the University of Paris IV-Sorbonne. We would like to thank Inge Arnesen, Johan van Benthem, Denis Bonnay, Julien Dutant, Pascal Engel, Krista Lawlor, Philippe Schlenker, Josh Snyder, Timothy Williamson as well as our students for most helpful discussion on this and related matters. We are also grateful to three anonymous referees for comments that led to a much improved version of the paper.

References

- Bonnay D. & Égré P. (2006). A Non-Standard Semantics for Inexact Knowledge with Introspection. In R. Parikh and S. Artemov (eds), *Proceedings of the Workshop Rationality and Knowledge*, ESSLLI 2006, Malaga.
- Bonnay D. & Égré P. (forthcoming), Margins for Error in Context. Forthcoming in M. Garcia-Carpintero & M. Kölbel (eds), *Relative Truth*, Oxford University Press.
- Dretske, F. (1993). Conscious perception. *Mind*. Vol. 102: 263-283.
- Evans, G. (1982), *The Varieties of Reference*. (Oxford, Clarendon Press.)
- Gomez-Torrente, M. (1997). Two Problems for an Epistemicist View of Vagueness. (In E. Villanueva (ed.), *Philosophical Issues*, 8: Truth, Ridgeview, Atascadero, 237-245.)
- Graff, D. (2002). An Anti-Epistemicist Consequence of Margin for Error Semantics for Knowledge. *Philosophy and Phenomenological Research* 64: 127-42.
- Gordon, R. (1995). Simulation without introspection or inference from me to you. (In M. Davies & T. Stone (eds), *Mental Simulation*, Oxford, Blackwell.)
- Hintikka, J. (1962). *Knowledge and Belief*. (Cornell University Press.)

Lewis D. (1974). *Intensional Logics without Iterative Axioms*. (Repr. in *Papers in Philosophical Logic*, Cambridge 1998: 66-76.)

Williamson, T. (1990). *Identity and Discrimination*. (Oxford, Blackwell.)

Williamson, T. (1992). Inexact Knowledge, *Mind* 101: 217-42.

Williamson, T. (1994). *Vagueness*. (London, Routledge.)

Williamson, T. (2000). *Knowledge and its Limits* (Oxford, Oxford University Press.)

ⁱ Williamson's first formulation of the margin for error principle appears in Williamson (1990: 104-106). His criticism of the (KK) principle originally appears in Williamson (1992). The argument is also presented in chapter 8 of Williamson (1994).

ⁱⁱ Following Lewis (1974), we call *iterative* a formula in which "intensional operators occur within the scope of intensional operators".

ⁱⁱⁱ We say that an iterative formula is *pure* if it involves only one kind of operator. Axiom (KF) is iterative and pure in this sense, whereas (KF'), although iterative, is not pure, since it involves two distinct modalities.

^{iv} Let's call KK' the bimodal system with two modalities K and K_π , closed under the two rules of generalization $\phi \vdash K\phi$ and $\phi \vdash K_\pi\phi$, and such that: K satisfies the schemata **4**, **T** and **K**; K_π satisfies the schemata **T** and **K**; axiom (KK') holds: $K_\pi\phi \rightarrow KK_\pi\phi$. The schemata $K(K_\pi\phi \rightarrow \phi)$ and $K_\pi\phi \rightarrow K\phi$ are theorems in KK'. The converse $K\phi \rightarrow K_\pi\phi$ is not a theorem of KK', however, for KK', with its usual Kripke semantics, is sound with respect to the class of frames (W, R, R_π) in which R is reflexive and transitive, R_π is reflexive, and $\forall xyz(xRy \ \& \ yR_\pi z \rightarrow xR_\pi z)$; it is easy to define a two-world model satisfying this condition in which $K\phi \rightarrow K_\pi\phi$ is not valid (the same result holds if one assumes K and K_π are both S5 modalities). Note that $R \subseteq R_\pi$ in such frames. KK' is already too strong a system for our purposes, however, since it also proves $K_\pi(K\phi \rightarrow \phi)$, by generalization, which would mean that I can perceive that my general knowledge is factive, something we want to avoid. What matters here, however, is that even in such a strong system there is no collapse between the modalities K and K_π .

^v We are indebted to an anonymous reviewer for this example.

^{vi} We could make that assumption without affecting our argument, however. Such an assumption is for instance present in Bonnay & Égré's (2006) epistemic semantics, in which only one epistemic operator is used, but such

that margins of error can be confined to first-order knowledge. Their approach makes clearer the fact that if I know visually that p , for instance, assuming this knowledge is constrained by a margin of error, and if q follows from p , then my first-order knowledge that q obeys the same margin of error, be it purely inferential or a mix of vision and inference. The important point, however, here as in their account, is that my *knowing that I know* is not necessarily subject to this margin of error.

^{vii} We do not thereby claim that any specific form of knowledge fails to be iterative as we claim of perceptual knowledge. For instance, although “I know by testimony that I know by testimony that p ” sounds odd (if we understand this as *de se* knowledge), we can perfectly imagine adverbial modifications other than “reflectively”, like “I know a priori that I know a priori that p ”, for which the iterations make sense. It remains possible to interpret “I know by testimony that I know by testimony that p ”, however, if the subject means that she *cannot recall* that *it is by testimony* that she came to know p , although someone reminded her that her knowledge that p comes from a specific testimony. If so, this means there is an interesting difference between knowledge by perception and knowledge by testimony: in the case of knowledge by perception, we claim that iterations of K_{π} are most likely ill-formed; by contrast, if there was an operator K_{τ} for knowledge by testimony, iterations would make sense in particular cases, but this would not be sufficient to support the validity of a principle of the form $K_{\tau}p \rightarrow K_{\tau}K_{\tau}p$.

^{viii} Note that we are *not* saying that Williamson himself would claim that all forms of knowledge are subject to the margin for error principle. Even though Williamson does not discuss the matter explicitly, the quote given in section 1 above suggests the opposite. Our claim is only that, in his argument against the (KK) principle, Williamson fails to take into account that different forms of knowledge are involved.

^{ix} As Evans (1982: 225-6) puts it, “whenever you are in a position to assert that p , you are *ipso facto* in a position to assert ‘I believe that p ’.”

^x Here, we adapt another formulation used by Williamson to define the notion of a margin for error principle: “A *margin for error principle* is a principle of the form ‘A’ is true in all cases similar to cases in which ‘It is known that A’ is true” (1994: 227).

^{xi} The armadillo example is borrowed from Dretske (1993), who argues that one can see an armadillo without seeing *that* it is an armadillo.

^{xii} A different line of criticism against the validity of margin for error principles has been investigated by M. Gomez-Torrente and D. Graff. The criticism is based on the idea that *some* propositions seem obviously known and also subject to the (KK) principle. Thus “we *know* that any man with zero hairs is bald”

(Graff 2002: 127), and all iterations of knowledge seem to hold as well. If the existence of such propositions is granted, and the margin for error principle is assumed to hold for all numerical values, the puzzle of Williamson that we examined comes up again.

^{xiii} A detailed discussion of the plausibility of (KME) is beyond the scope of its paper. For a precise discussion of the principle, see Bonnay & Egré (forthcoming), who make an explicit comparison with the present approach.