
Inexact Knowledge, Margin for Error and Positive Introspection

Julien Dutant

Philosophy Dept., University of Geneva
Institut Jean Nicod, Paris
Department of Philosophy,
University of Geneva,
rue de Candolle, 2
CH-1211 Geneva 4, Switzerland

Abstract

Williamson (2000a) has argued that positive introspection is incompatible with inexact knowledge. His argument relies on a margin-for-error requirement for inexact knowledge based on a intuitive safety principle for knowledge, but leads to the counter-intuitive conclusion that no possible creature could have both inexact knowledge and positive introspection. Following Halpern (2004) I put forward an alternative margin-for-error requirement that preserves the safety requirement while blocking Williamson’s argument. I argue that the infallibilist conception of knowledge that underlies the new requirement provides a better account of inexact knowledge and higher-order knowledge than both Williamson’s and Halpern’s.

The positive introspection axiom (or principle KK) states that if one knows, one knows that one knows (Hintikka, 1962). It is routinely used in formal epistemology, at least as a reasonable idealisation. Inexact knowledge is knowledge that fails to distinguish close possibilities: for instance, one might know that a pen is not 3 inches long without knowing whether it is 2 inches or 2.1 inches long. Clearly, most of our knowledge is inexact knowledge. However, Williamson (1992, 1994, 2000a) has argued that positive introspection is incompatible with a plausible margin-for-error requirement for inexact knowledge. The requirement is based on the appealing claim that knowledge requires one’s belief to be safely true. But if Williamson is right, it leads to the unexpectedly strong conclusion it is impossible for any agent, however idealised, to reach arbitrarily high orders of knowledge with respect to any item of inexact knowledge.

Several authors have responded by rejecting the safety requirement altogether. By contrast, I will for-

mulate an alternative margin-for-error requirement that both preserves the safety intuitions and blocks Williamson’s argument against positive introspection. The proposed model of inexact knowledge is based on (Halpern, 2004), but unlike Halpern’s, it does not validate positive introspection either. Instead, it provides conditions at which positive introspection may hold. I will argue that the resulting view provides a better account of inexact knowledge and higher-order knowledge than alternative accounts.

Williamson’s argument is restated in section (1). The alternative margin-for-error requirement is formulated in section (2). Section (3) extends the account to higher-order knowledge and provides motivations to prefer it over Williamson’s and Halpern’s.

1 Williamson’s anti-luminosity argument

In this section, I will state the argument, review available responses, and formulate a semantics for “know” that validates Williamson’s rejection of (KK).

1.1 Williamson’s argument

Williamson (2000a) relies on the following situation of inexact knowledge.¹ Mr Magoo wonders how tall a certain tree is. Magoo’s estimations of the tree’s height are imperfect, but they do allow him to gain some knowledge of the tree’s height, for instance that it is not 10 nor 1000 inches high. Let us assume that Magoo is liable to make errors of plus or minus 1 inch about the tree’s height. For simplicity, I will abbreviate “the tree is i inches tall” by “the tree is i ”. Thus, for no i , Magoo knows that the tree is i . However, for some i ,

¹In Williamson’s terminology, a condition (such as one feeling cold or one’s knowing p) is *luminous iff* if the condition holds, one is in position to know that it holds. Thus Williamson’s argument against positive introspection is an “anti-luminosity” argument.

Magoo knows that the tree is not i .

According to Williamson (2000a), Magoo cannot know of a tree of i that it is not $i + 1$. Even if he correctly judged that the tree was not $i + 1$, he would not be safe from error, because he could have easily made the same judgment about a tree of $i + 1$. Thus Magoo's knowledge is subject to the following Margin-for-Error requirement:

(ME_W) For all i , if the tree is $i + 1$, then Magoo does not know that it is not i .

Let p_i be the proposition that the tree is i inches tall, Kp the proposition that Magoo knows p , and let " \rightarrow " denote material implication:

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

Williamson shows that if Magoo knows (ME_W), as seems possible, positive introspection leads to a contradiction. Let the tree be 100. Magoo knows that it is not 0. By introspection, he knows that he knows it. Since he knows (ME_W), he knows that, if he knows it, the tree is not 1. By inference, he knows that the tree is not 1. By reiterating those steps, he knows that the tree is not 100. But that is impossible, since knowledge entails truth and the tree is 100.

More rigorously, we need four background assumptions. First, the tree has some height, say 100, (F) and second, Magoo has some inexact knowledge of it, namely that it is not 0 (ik). Third, knowledge is closed under known implication (EC). Fourth, knowledge is factive (T):

(F) p_{100}
 (ik) $K\neg p_0$
 (EC) for any $p, q: (Kp \wedge K(p \rightarrow q)) \rightarrow Kq$
 (T) for any $p: Kp \rightarrow p$

The main premisses are Magoo's knowledge of (KME_W) and positive introspection of his knowledge of the tree's height:

(KME_W) $\forall i K(p_{i+1} \rightarrow \neg K\neg p_i)$
 (KK) $\forall i (Kp_i \rightarrow \neg K\neg p_i)$

The argument goes as follows. Suppose the tree is 100:

1	$K\neg p_0$	(ik)
2	$KK\neg p_0$	by (1), (KK)
3	$K(K\neg p_0 \rightarrow \neg p_1)$	by (KME _W)
4	$K\neg p_1$	by (2), (3), (EC)
...	(same steps for heights 1, 2, ...)	
n	$K\neg p_{100}$	
	$\neg p_{100}$	by (n), (KF)
	\perp	

Two remarks. First, the situation involves several simplifications: Magoo's inaccuracy is itself precise and constant across heights, and he knows it accurately.² However, the argument would go through if we assumed instead that Magoo's inaccuracy was always substantially more than a tenth of an inch, and that Magoo knew that it is at least a tenth of an inch. Second, as Williamson (2000a, p.118-9) argues, the argument does not rely on vagueness in the verb "knows". In a standard sorites argument, the induction premisses (*e.g.* $Bald(i) \rightarrow Bald(i + 1)$) is intuitively justified because of the predicate's vagueness. But in the present argument, the corresponding claim ($K\neg p_i \rightarrow \neg p_{i+1}$) is derived, and the argument goes through even if "knows" is sharpened.

1.2 Responses to the argument

One of the premisses or assumptions must go. I will not consider the extreme options of rejecting facts (F), knowledge (ik) (scepticism), factivity (KF). Epistemic Closure (EC) is much discussed, but as Williamson (2000a, p.116-8) notes, the argument requires only a restricted version that appears harmless.

Williamson rejects (KK), which seems reasonable since (KK) is implausible anyway. However, giving up (KK) *because of the anti-luminosity argument* has problematic consequences. First, even more plausible versions of (KK) are ruled out, *e.g.* if one knows, one is in position to know that one knows. Second, no possible agent, however idealised, could satisfy (KK) with respect to any item of inexact knowledge.³ Third, every iteration of knowledge is costly: Magoo knows that the tree is not i only if the tree is $i + 1$ or more, but he knowsⁿ that he knows it only if it is $i + n + 1$ or more. Yet we find it hard to make any difference between higher orders. Fourth, no item of inexact knowledge can be known at arbitrarily high orders. That is surprising for Magoo's knowledge that the tree is not 0.001 inch.

Several authors have rejected (ME_W) by arguing

²see Williamson (2000b)

³Unless one preserves (KK) by making *ad hoc* move of restricting it to some subset of the propositions about the tree's height.

against the safety requirement on which it is based.⁴ However, as Williamson (2005) notes, their putative counter-examples to a generalised safety principle have no bearing on the limited margin-for-error requirements used in the argument. Moreover, rejecting one instance of (ME_W) seems incompatible with the idea that Magoo’s knowledge is inexact: for instance, Magoo would be able to know that the tree is not 100 when it is just 99.

Bonnay and Egré (2006) reject (KME_W) without rejecting (ME_W). But their view has the implausible consequence that Magoo could not even know a weaker version of his actual margin-for-error requirement – say, if he takes the margin to be 0.1 at least while it is 1)

Dokic and Egré (2004) and Egré (2007) advocate a different solution, by distinguishing two types of knowledge: perceptual knowledge that satisfies (ME_W) but not (KK); reflexive knowledge that satisfies (KK) but not (ME_W). The view has to treat “know” as ambiguous and implies that margins are not iterated at all, which has implausible consequences (see section 3.3).

My proposal is to reject (ME_W) by providing an alternative margin-for-error requirement that preserves the safety intuitions. Similar proposals have been hinted at by Mott (1998) and made by Halpern (2004).

1.3 Williamson’s intransitive model of inexact knowledge

Williamson (1992) formulates a simple propositional modal logic KT for knowledge that validates (ME_W). The logic crucially involves an intransitive accessibility relation.

Let $M = \langle W, \sim, V \rangle$ be a Kripke structure where $W = \{w_1, w_2, \dots, w_n\}$ ($n \in \mathbb{N}$), and $w_i \sim w_j$ iff $|i - j| \leq 1$. V is an interpretation function such that $V(p_i) = w_i$ for $0 \leq i \leq n$. Intuitively, the possible worlds w_i correspond to the possible heights of the tree, each p_i is true exactly at the corresponding w_i and the relation $w \sim w'$ represents Magoo’s inability to discriminate between trees whose height differ by 1 or less. It is easy to check that the relation is reflexive, symmetric but not transitive ($1 \sim 2$ and $2 \sim 3$ but not $1 \sim 3$). We define satisfaction as follows:

- $(M, w) \models p_i$ iff $w \in V(p_i)$
- $(M, w) \models \neg p$ iff it is not the case that $(M, w) \models p$
- $(M, w) \models (p \wedge q)$ iff $(M, w) \models p$ and $(M, w) \models q$
- $(M, w) \models Kp$ iff for all w' s.t. $w' \sim w$, $(M, w') \models p$

The model validates (ME_W): if $K\neg p_{100}$ is true at w_k

⁴Brueckner and Fiocco (2002), Conee (2005), Neta and Rohrbaugh (2004).

then $\neg p_{100}$ is true at w_{k+1} , so w_{k+1} cannot be a world in which the tree is 100, and thus w cannot be a world where the tree is 99. More generally, for any i , if $K\neg p_i$ then $\neg p_{i-1}$.

For knowledge to be iterated, the semantics requires p to be true at further worlds: Kp is true at w_i iff for any j s.t. $|i - j| \leq 1$, p is true at w_j , but KKp is true at w_i iff for any j s.t. $|i - j| \leq 2$, p is true at w_j . Each iteration requires a wider margin for error, and that is why (KK) cannot hold for any contingent proposition.

2 An alternative model for inexact knowledge

Williamson (2000a, p.115) introduces (ME_W) in the following passage:

To know that the tree is i inches tall, Mr Magoo would have to judge that it is i inches tall; but even if he so judges and in fact the tree is i inches tall, he is merely guessing; for all he knows it is really $i - 1$ or $i + 1$ inches tall. He does not know that it is not. Equally, if the tree is $i - 1$ or $i + 1$ inches tall, he does not know that it is not i inches tall.

The last sentence states (ME_W). But the first states a different requirement, which I shall now present.

2.1 Halpern’s framework and the revised margin-for-error principle

In Williamson’s framework, worlds are individuated by objective facts such as the tree’s being 100. Halpern (2004) advocates an alternative framework in which worlds are individuated by objective facts *and the subject’s state*. For instance, given Magoo’s limited powers of discrimination, there are three possible cases in which the tree is 100:

1. The tree is 100 and Magoo’s estimation is 99.
2. The tree is 100 and Magoo’s estimation is 100.
3. The tree is 100 and Magoo’s estimation is 101.

Conversely, if Magoo estimates that the tree is 100, it might really be 99, 100 or 101. We thus have a matrix of worlds:

E101		w	w
E100		w	w
E99	w	w	w
E98	w	w	
E97	w		
	98	99	100
		101	

The possible situations are marked as “w”. The situation (100, E101), for instance, is the situation in which the tree is 100 Magoo estimates it as being 101.

Suppose that Magoo estimates that the tree is 100. If he judged that it was not 99, he would not be safe from error, because he might have made the same estimation if the tree was 99, as (99, E100) shows. Thus he cannot know that the tree is not 99. That suggests a revised Margin-for-Error requirement:⁵

(ME_R) For all i , if Magoo estimates that the tree is $i + 1$, then Magoo does not know that it is not i .

Let Ep_{*i*} be the proposition that Magoo estimates that the tree is i :

(ME_R) $\forall i (Ep_{i+1} \rightarrow \neg K\neg p_i)$

The condition is sufficient to rule out Mr. Magoo’s having exact knowledge. 1) For no size i does Magoo know that the tree is i . Even if he estimates that the tree is i , he might have made the same estimation if the tree was not i . 2) For any size i it is possible that the tree is not i and Magoo fails to know that it is not i . That might happen if the tree is $i + 1$ or $i - 1$.

In some respects, (ME_R) is stronger than (ME_W). At (99, E100), (ME_W) does not rule out Magoo’s knowing that it is not 101. But (ME_R) does so, because Magoo could have made the same estimation if the tree was 101. In other respects, (ME_R) is weaker. At (100, E99), (ME_W) rules out Magoo’s knowing that it is not 101. But (ME_R) does not so, because Magoo would not have made that estimation if the tree was 101.

To see what goes on, consider the following conditionals:

- (a) If the tree is 100, Magoo knows that it is not 99.
- (b) If the tree is 100, Magoo does not know that it is not 99.

Williamson’s framework makes it seem that either (a) or (b) is true. (In the corresponding model, that is reflected in the fact that there is only one possible world per tree height.) Since (a) grants Magoo with exact knowledge, Williamson holds (b) – and more generally (ME_W). But (ME_R) implies only that if the tree is 100, then Magoo *might* fail to know that it is not 99. Conversely, (ME_R) does not rule out that *Ma-*

⁵Mott (1998) has suggested a similar requirement. Halpern (2004)’s “reports” operator is subject to the same requirement.

goo knows that the tree is not 99 and the tree is 100 (100, E99).

Now, one can see why Williamson’s argument cannot go through with the new requirement alone. Suppose Magoo knows that the tree is not 99. (ME_R) does not rule out that the tree is 100. So even if Magoo knows that he knows it, he cannot infer that the tree is 100. The sorites-like series cannot be initiated.

The contrast between (ME_W) and (ME_R) is shown below. Let p be the proposition that the tree is not 101. By Williamson’s requirement, Magoo fails to know p if the tree is 101 ± 1 :

E101			$\neg Kp$	$\neg Kp$
E100	w		$\neg Kp$	$\neg Kp$
E99	w	w	$\neg Kp$	
E98	w	w		
E97	w			
	98	99	100	101

The revised requirement rules out situations in which Magoo *estimates* the tree as being 101 ± 1 :

E101			$\neg Kp$	$\neg Kp$
E100		$\neg Kp$	$\neg Kp$	$\neg Kp$
E99	w	w	w	
E98	w	w		
E97	w			
	98	99	100	101

To sum up, there are two margin-for-error requirements that can be claimed to capture the limitations of inexact knowledge. They are merely necessary conditions on knowledge, and they are compatible though distinct. But for his anti-luminosity argument Williamson needs to argue that inexact knowledge is subject to his requirement on top of, or instead of, the revised one.

2.2 Halpern’s transitive model for inexact knowledge

Following Halpern (2004), we can formulate a S5 modal logic in which (ME_R) and (KK) hold.

Let $M = \langle W, \sim, V \rangle$ be a Kripke structure where $W = \{w_{i,j} | i \in \mathbb{N}, j \in \mathbb{N}, |i - j| \leq 1\}$, $w_{i,j} \sim w_{i',j'}$ iff $j' = j$, and V is such that $V(p_i) = \{w_{i',j'} | i' = i\}$. Intuitively, $w_{i,j}$ is the world in which the tree is i and Magoo estimates it as being j ; given Magoo’s powers of discrimination, there is no world in which he makes a mistake of more than one inch. The accessibility relation relates two worlds only if Magoo makes the same estimation in both. We define satisfaction as before.

The model validates (ME_R): if Magoo estimates that the tree is k , then there is an accessible world (namely, $(k-1, Ek)$) in which the tree is $k-1$, so Magoo cannot know that the tree is not $k-1$.

The model validates (KK): the accessibility relation is transitive, since each world has access to all and only worlds in which the estimation is the same. From this (KK) follows straightforwardly.

Interestingly, the model invalidates (ME_W). At (100, E99), Magoo knows that the tree is not 101, since it is not 101 in all accessible worlds: (98, E99), (99, E99), (100, E99). So $Kp_i \rightarrow \neg p_{i+1}$ is not valid.

3 Extension to higher-order knowledge

3.1 The underlying infallibilist conception of knowledge

The intuition behind the revised margin-for-error principle is *infallibilism*: Magoo knows that p only if he could not have been in his actual subjective state if p was false. That can be turned into a *sufficient* condition for knowledge. Let E_p be Magoo's p -relevant estimation:

$$(IK) \quad Kp \text{ iff } \Box(E_p \rightarrow p)$$

The intuition behind (IK) is the following. Suppose Magoo estimates that the tree is 100. Given his powers of discrimination, he could not have made that estimation if the tree was 98. So he could not have been wrong about the tree not being 98. That seems sufficient for him to know that it is not 98. (Halpern's model validates (IK); the accessibility relation for \Box is just $W \times W$. More generally, (IK) implies that Williamson's (ME_W) is false.)

Infallibilism raises sceptical worries: surely, a mad scientist could have Magoo estimating that the tree is 100 while there is no tree. However, that can be avoided in several ways, for instance by restricting the necessity operator in the contextualist's way (remote possibilities are ignored in common contexts) or the subject-sensitive invariantist's way (remote possibilities are not genuine in S 's particular situation). The details do not matter here.

Estimations are subjective states that S could be in only if certain facts obtain. They roughly correspond to Lewis (1996)'s notion of evidence. Thus when I say that an estimation E is incompatible with p , I do not mean that E 's content implies not- p , but that S 's *being in* E implies not- p . I need not specify them further

here: they can be experiences, judgements about the way things look, or beliefs.

3.2 Extension to higher orders

The natural way of extending (IK) to higher orders is to introduce estimations of estimations:

Magoo knows that he knows that the tree is not 98 *iff* necessarily, if he estimates that he estimates that the tree is 100, he knows that the tree is not 98.

Again, estimations of estimations are just states S could be in only if certain facts (here, estimations) obtain. They can be judgments about one's experiences, beliefs caused by those experiences, or more generally one system's information about another system's informational state, whether in the brain or in an artificial system.⁶

Thus, let E^2p be Magoo's estimation of his p -relevant estimation, and similarly for $E^n p$:

$$(IK^n) \quad K^n p \text{ iff } \Box(E^n p \rightarrow p)$$

The left-to-right part of (IKⁿ) implies that an equivalent of the revised margin-for-error requirement is true at each level.⁷

Consequently, (KK) holds between two orders only if Magoo's higher-order estimation of his lower-order state is perfect. Assume Magoo is liable to make errors of 1 about his lower-order estimations – which is likely if they are experiences. Suppose he estimates the tree as being 100, so he knows that it is not 98. He also estimates rightly that he estimates so, but that is compatible with his estimating the tree as being 99, in which case he would not have known that the tree is not 98. So he does not know that he knows it. Thus (IKⁿ) is compatible with (KK) but does not entail it.

Formally, we need a yet richer set of worlds: if Magoo is liable to such higher-order errors, we need to distinguish (100, E100, EE100) from (100, E100, EE99) and (100, E100, EE101). Accessibility on these new worlds is such that $(i, E_j, EE_k) (i', E_{j'}, EE_{k'})$ iff $k = k'$. On this view, Halpern's model represents the degenerated case in which there is only one EE world for each E world: that is, all Magoo's higher-order estimations are exact. That is why the model validates KK.

⁶Thanks to an anonymous referee here.

⁷To avoid inconsistency, we assume that higher-order estimations are not more sensitive to facts than lower-order ones.

3.3 Defence of the revised view

There are three reasons to adopt the revised view of inexact and higher-order knowledge presented here:

First, Williamson's (ME_W) implies that inexact knowledge states are strictly more informative than their content. For instance, if Magoo knows that the tree is between 99 and 101, then the tree is neither 99 nor 101. So the information carried by Magoo's knowledge state is richer than what Magoo knows. That is the deep reason for (KK)'s failure: additional information is required for Magoo to know that he knows it. On the revised conception, *knowledge states are just as informative as their content*: Magoo's knowing that the tree is between 99 and 101 only indicates that the tree is between 99 and 101. That is welcome, since there is no reason why the information carried by a subject's state should necessarily be partly hidden to her.

Second, (ME_W) implies that subjects infer their own states from the state of the world, in spite of direct information they might have about them. Consider why (ME_W) rules out Magoo's knowing that he knows that the tree is not 101 at (99, E99). For all Magoo knows, the tree might be 100. But if it was 100, Magoo might be *estimating* it as being 100. Hence, the reasoning goes, for all Magoo knows, he might be estimating it as being 100. (Therefore, for all he knows, it might be 101.) But that move is unacceptable if Magoo is aware that he does *not* estimate it as being 100. So (ME_W) is cogent only if Magoo has no direct information about his own states. By contrast, on the revised view, the epistemically accessible worlds are only worlds in which Magoo makes the same estimation. For instance, if Magoo estimates that he estimates the tree as being 99, and would not do so if he estimated it as being 101, then it is false that for all he knows, he might be estimating it as being 101. Thus the revised view *allows subjects to have direct information about their own states*.

Third, Halpern's view implies that subjects cannot *lack* information about their own states.⁸ Suppose Magoo estimates wrongly that he estimates the tree as being 101 – in fact, he estimates it as being 100. On Halpern's view, that is sufficient for him to know that he knows that the tree is not 101.⁹ On the present

⁸The same holds for other models that validate (KK), such as Bonnay and Egré (2006) or Egré (2007).

⁹Note that the fact that Halpern's model validates (KK) is not simply a consequence of the accessibility relation's being Euclidean and transitive. In the model presented here, the accessibility relation is such, and yet (KK) need not hold. That is because in those models, whether (KK) holds depends on whether there are several dimensions of subjective states. (Thanks to an anonymous referee here.)

view, it is not, because Magoo could have made the same second-order estimation if he did *not* know that the tree was not 101. More generally, (IK^n) is compatible with different margins for error at each order, or the absence of them. This is welcome, since subjects *can have merely partial information about their own states*.

The revised view thus leaves it open whether (KK) holds between any two orders. For instance, it is compatible with the following attractive view: from a certain order n on, the very same cognitive state of ours implements both our estimation ^{n} that p , our estimation ^{$n+1$} that p and so on. That is reflected in our ability to distinguish them. But by the same token, we cannot fail to be estimating ^{n} that p when we estimate ^{$n+1$} that p , so (KK) holds from n on. By contrast, we do distinguish the first orders but are liable to make errors about our own states at those orders, so (KK) fails to hold there. On that view, if the tree is 100, the closer i is to 100, the lesser Magoo is likely to know that he knows that the tree is not i ; nevertheless, he has arbitrary high orders of knowledge that it is not 0.

4 Conclusion

I have defended a revision of Williamson (2000a) Margin-for-Error principle for inexact knowledge that preserves the safety intuitions while avoiding the implausible consequence that no possible creature can satisfy (KK) with respect to any piece of inexact knowledge. On the basis of the revised principle, I have defended an infallibilist conception of knowledge that provides a more flexible account of higher-order knowledge, allowing (KK) to hold or not between any two orders, and giving the condition at which it does. I have argued that the resulting view also avoids two problematic consequences of Williamson's account: namely, that some information carried by knowledge states is necessarily hidden, and that subjects cannot rely on direct information about their own states, and a problematic consequence of Halpern's account: that subjects cannot lack information about their own states.

Acknowledgements

The present paper has been inspired by the work by Paul Egré, Jérôme Dokic and Denis Bonnay on these issues. Many thanks to them and to Ghislain Guigon for useful discussions. Thanks also to Pascal Engel, Gianfranco Soldati, Philipp Keller, Eric Olsson, and the audience in Geneva for their questions and comments; and to an anonymous referee for helpful suggestions.

References

- Denis Bonnay and Paul Egré. A non-standard semantics for inexact knowledge with introspection. In Sergei Artemov and Rohit Parikh, editors, *Proceedings of the ESSLLI 2006 Workshop on Rationality and Knowledge*. Malaga, 2006.
- Anthony Brueckner and Mark Oreste Fiocco. Williamson's anti-luminosity argument. *Philosophical Studies*, 110:285–293, 2002.
- Earl Conee. The comforts of home. *Philosophy and Phenomenological Research*, 70(2):444–51, 2005.
- Jérôme Dokic and Paul Egré. Margin for error and the transparency of knowledge. Technical report, Institut Jean Nicod, 2004. URL http://jeannicod.ccsd.cnrs.fr/ijn_00000479/en/. Submitted.
- Paul Egré. Reliability, margin for error and self-knowledge. In Vincent Hendricks and Duncan Pritchard, editors, *New Waves in Epistemology*. Ashgate, 2007.
- Joseph Y. Halpern. Intransitivity and vagueness. In Didier Dubois, Christopher Welty, and Mary-Anne Williams, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Ninth International Conference (KR2004)*. The AAAI Press, Menlo Park, California, 2004.
- Jaakko Hintikka. *Knowledge and Belief*. Cornell University Press, 1962.
- David K. Lewis. Elusive knowledge. *Australasian Journal of Philosophy*, 74:549–567, 1996.
- Peter Mott. Margins for error and the sorites paradox. *Philosophical Quarterly*, 48(193):494–504, 1998.
- Ram Neta and Guy Rohrbaugh. Luminosity and the safety of knowledge. *Pacific Philosophical Quarterly*, 85(4):396–406, 2004.
- Timothy Williamson. Inexact knowledge. *Mind*, 101: 218–42, 1992.
- Timothy Williamson. *Vagueness*. Routledge, London, 1994.
- Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, 2000a.
- Timothy Williamson. Margin for errors: A reply. *Philosophical Quarterly*, 50(198):76–81, 2000b.
- Timothy Williamson. Replies to commentators. *Philosophy and Phenomenological Research*, 70(2):468–91, 2005. to appear.