

## Epistemic closure under deductive inference: what is it and can we afford it?

Assaf Sharon · Levi Spectre

Received: 4 January 2011 / Accepted: 3 February 2012 / Published online: 28 February 2012  
© Springer Science+Business Media B.V. 2012

**Abstract** The idea that knowledge can be extended by inference from what is known seems highly plausible. Yet, as shown by familiar preface paradox and lottery-type cases, the possibility of aggregating uncertainty casts doubt on its tenability. We show that these considerations go much further than previously recognized and significantly restrict the kinds of closure ordinary theories of knowledge can endorse. Meeting the challenge of uncertainty aggregation requires either the restriction of knowledge-extending inferences to single premises, or eliminating epistemic uncertainty in known premises. The first strategy, while effective, retains little of the original idea—conclusions even of modus ponens inferences from known premises are not always known. We then look at the second strategy, inspecting the most elaborate and promising attempt to secure the epistemic role of basic inferences, namely Timothy Williamson’s safety theory of knowledge. We argue that while it indeed has the merit of allowing basic inferences such as modus ponens to extend knowledge, Williamson’s theory faces formidable difficulties. These difficulties, moreover, arise from the very feature responsible for its virtue- the infallibilism of knowledge.

**Keywords** Knowledge · Inference · Modus ponens · Infallibilism · Lottery propositions · Chance · Epistemic probability · Knowledge safety · Single premise closure · Multi premise closure

---

A. Sharon  
Stanford University, Stanford, CA, USA  
e-mail: Assafsh@stanford.edu

L. Spectre (✉)  
The Open University of Israel, Raanana, Israel  
e-mail: Levi.spectre@gmail.com

## 1 knowledge by inference

I believe many things, but I also know I'm wrong about some of them. There is nothing irrational, bad, or paradoxical about my epistemic situation. In fact, it seems that there's something wrong with subjects who don't know, or at least strongly believe, that they are mistaken about one or more of their beliefs. This mundane and seemingly innocent observation, nonetheless, is of great consequence for the question of whether the sets of propositions that are believed or known are closed under certain logical operations. Standardly understood, a set is closed under a logical operation if and only if the result of the logical operation will be a member of the set. Specifically, since belief and knowledge are closely connected to truth and since the set of all true propositions is surely closed under logical operations that preserve truth, knowledge and belief will be closed under any operation that is *salve veritate*. Thus, the set of all known propositions would be closed under deduction. The observation above, however, shows that this is not the case.<sup>1</sup> Although the truth of propositions believed or known is certainly closed under any truth-preserving operation, the attitudes are not.<sup>2</sup> Preface-type considerations show that the transition from a conjunction of many items of belief to a belief in the conjunction of those items, is not innocent. If I believe each of  $p_1, \dots, p_n$ , then assuming  $n$  is sufficiently high, I ought not believe the conjunction  $p_1 \wedge \dots \wedge p_n$ . This is due to the by now familiar phenomenon of accumulation of uncertainty (or risk of error).

Still, it seems intuitively appealing to maintain that even if belief is not closed in the strict technical sense explicated above, it is quasi-closed. Thus, for example, if the set of propositions believed by a rational agent includes  $p$  and includes *if  $p$  then  $q$* , then the set should also include  $q$  (or be revised to exclude one of the former beliefs). Clearly, this would be giving up the thrust of closure, namely the idea that these attitudes are closed under truth preserving operations. Nevertheless, if a more restricted operation can be delineated under which attitudes such as knowledge and belief are closed, this could be of great significance. Alas this too will not work, for the same threat of uncertainty accumulation arises here as well. Suppose I believe that  $p$  to degree 0.7 and that *if  $p$  then  $q$* , to the same degree. My rational degree of belief in  $q$  (at least for some cases)<sup>3</sup> should be 0.49. And so although I ought not believe in  $q$  I need not disbelieve that  $p$  or that *if  $p$  then  $q$* .<sup>4</sup>

Despite the familiarity of these considerations regarding belief, it is widely accepted that knowledge behaves differently. Specifically, most theorists take it to be obvious that knowledge is closed under logical inference, or at least that it is closed under

<sup>1</sup> For the original arguments for this claim see Kyburg (1961, 1970) and Makinson (1965).

<sup>2</sup> Even if psychological barriers such as lack of formation of the relevant belief are overcome.

<sup>3</sup> If the rational probability function for a subject is  $\Pr(\bullet)$ , then given that  $\Pr(p) = 0.7$  and  $\Pr(p \rightarrow q) = 0.7$ , then (assuming independence)  $\Pr(q)$  may be 0.49. That is, if the rational probability of the conjunction of the premises— $\Pr(p \wedge (p \rightarrow q))$ —is 0.49.

<sup>4</sup> Christensen (2004) makes this argument with respect to beliefs understood as conforming to the probability calculus. The argument here does not make such an assumption—as long as uncertainty is allowed to accumulate, the problem will arise whether or not beliefs are gradable. This allows us to apply the argument to knowledge, which remains unaddressed in Christensen's book.

competent deduction. That is, if one *knows* that  $p$  and *knows* that *if  $p$  then  $q$* , and one infers  $q$  competently from these premises while retaining knowledge of the premises throughout, then one *knows* that  $q$ . In fact, this is precisely the way many theorists formulate the principle of epistemic closure. Moreover, it is the idea that knowledge can be extended by basic inferences such as modus ponens expressed in this formula that is often proposed as the main motivation for accepting the principle of epistemic closure. As basic and compelling as this idea may appear, it is, however, *un* acceptable, at least for standard accounts of knowledge.

It is easily shown that so long as knowledge is compatible with some degree of rational uncertainty, preface type considerations apply to knowledge just as they do to belief. One can know each of  $p_1, \dots, p_n$ , yet fail to know, e.g. since it is not even rational to believe, that  $p_1 \wedge \dots \wedge p_n$ . Yet, since knowledge presumably requires high degree of rational belief, it might seem that if one infers by modus ponens from known premises, one knows the conclusion. That is, necessarily if S knows that  $p$  and S knows that *if  $p$  then  $q$* , then S knows  $q$  (or is thereby in a position to know that  $q$ ). This, however, is mistaken as shown by the following argument:<sup>5</sup> Suppose I know that  $p$  and I know that  $q$  and I know that  $r$ .

- (i)  $Kp, Kq, Kr$   
I also know *a priori* that *if  $p$ , then (if  $q$ , then ( $p$ -and- $q$ ))*:
- (ii)  $K(p \rightarrow (q \rightarrow (p \wedge q)))$   
Now, since we are assuming that modus ponens is epistemically closed, I can infer and therefore know that *if  $q$ , then ( $p$ -and- $q$ )*.
- (iii)  $K(q \rightarrow (p \wedge q))$   
Since I know that  $q$  (premise i) a further modus ponens inference, will provide me with knowledge that ( $p$ -and- $q$ ).
- (iv)  $K(p \wedge q)$

Because I also know that  $r$ , I can rehearse this inference pattern to derive knowledge of  *$p$ -and- $q$ -and- $r$* . In fact I can in this way derive knowledge of the conjunction of everything I know.

Thus the phenomenon of accumulation of uncertainty gives us compelling reasons to reject the idea that knowledge transmits across even the most basic inferences such as *modus ponens*, *modus tolens*, etcetera. In other words, the rejection of multi-premise closure, a principle that guarantees knowledge of long conjunctions on the basis of knowledge of the conjuncts, is inconsistent with the endorsement of closure of knowledge under inference patterns such as modus ponens.

The argument above shows that belief and knowledge are closed under proper inference only if uncertainty is not allowed to accumulate. This can be achieved in one of two ways: First, if inferences are from one premise only (and the premise strictly implies the conclusion)<sup>6</sup>; second, if there is no uncertainty to begin with. The first of these options avoids the problem but only at the cost of making beliefs and items

<sup>5</sup> The argument is inspired by Peter Pagin's (1990) argument about rational belief.

<sup>6</sup> It is important to note that strict implication is not exclusive to logical reasoning. I cannot logically and validly derive my existence from the proposition that I am thinking, and yet, the kind of accumulation of doubt that we are appealing to will not mount in this case.

of knowledge inferentially isolated. This means that the epistemic status of inferred beliefs will only be guaranteed for those inferred *a priori* from single beliefs that enjoy this status. The original idea of a closed set is dramatically impoverished, obviously, by this restriction, for closure no longer pertains to subsets, but rather only to operations on single members of the set. Moreover, this option takes most of the bite out of the idea that inference is a proper way of expanding one's epistemic repertoire.<sup>7</sup> It leaves out most of the interesting and informative inferences we make in science as well as in our daily lives, inferences that contain combinations of items of knowledge or belief. If all we can salvage are strictly implied propositions from single items of knowledge, extension of our knowledge by inference is surely an impoverished notion. These are usually trivial, and when they appear interesting it is because they get epistemology into trouble (often skeptical trouble, e.g. "I know I have a hand" entails that I'm not a brain in a vat, etc.).<sup>8</sup>

Saving the epistemic role of basic inferences such as modus ponens is surely desirable, if it can be achieved. It is therefore of great interest to examine the success of the second, more ambitious strategy, namely eliminating the threat of uncertainty accumulation by eliminating uncertainty from the outset. Clearly to the extent that this has any prospects of success, it is only so for knowledge. This is because the idea that beliefs or even all rational beliefs are held with absolute certainty is implausible. In this paper we look at the most influential and elaborate attempt at pursuing this option, namely the safety theory of knowledge developed by Timothy Williamson. Saving the epistemic role of modus ponens inferences seems to us the best motivation for theories such as Williamson's, although, surprisingly, he does not make this connection explicitly. While it manages to rescue the epistemic role of basic inferences, we argue that it does so at significant, perhaps unbearable, cost.

## 2 Safety and chance

According to the safety theory of knowledge, beliefs that are safely true count as knowledge. Since the notion of "safely true belief" is closed under known implication, so is knowledge.<sup>9</sup> Focusing on this feature of the view, John Hawthorne and Maria Lasonen-Aarnio have recently made an important contribution to a more com-

<sup>7</sup> It is for this reason, we suppose, that many theorists, some even while repudiating multi-premise closure, formulate closure with two premises (necessarily, if  $(Kp \wedge K(p \rightarrow q))$ , then it follows that  $Kq$ ). Some even confusedly call it single-premise closure (Fumerton 2006, pp. 24–25). However, this formulation is in fact an instance of multi-premise closure as shown by the argument above. It would be true if the premises can always be known as a conjunction, i.e. if  $Kp \wedge K(p \rightarrow q)$  entails  $K(p \wedge (p \rightarrow q))$ . But this is to beg the question as the validity of multi-premise closure is already assumed. The relevant individuation conditions of 'premise' are determined by the range of the relevant operator (in this case knowledge) since the degree of uncertainty applies to the operator. Notable exceptions are Williamson (2000, p. 177) and Hawthorne (2004, pp. 32–34) who are careful to correctly distinguish single- from multi-premise closure. The crucial point we mean to stress is that the former does not entail closure of basic inferential modes such as modus ponens.

<sup>8</sup> We elaborate on this and related issues in our "Evidence and the Openness of Knowledge" *Philosophical Studies*, (forthcoming).

<sup>9</sup> Whether or not safety does in fact entail closure is an issue we will not discuss here. Our view is that without adding dubious assumptions regarding the basis for beliefs that count as knowledge, knowledge

prehensive understanding of the weakness of Williamson’s theory of knowledge. After presenting their insightful argument, we will assess Williamson’s response and show that it leads to further complications. Here is a formulation of the argument adapted from Williamson:

- (1)  $p_1, \dots, p_n$  are true propositions about the future.
- (2) Each of  $p_1, \dots, p_n$  has the same high chance (objective probability) less than 1.
- (3)  $p_1, \dots, p_n$  are probabilistically independent of each other (in the sense of chance).
- (4) The chance of  $p_1 \wedge \dots \wedge p_n$  is low [for large  $n$ ]. (from 2, 3, 4)
- (5) One believes  $p_1 \wedge \dots \wedge p_n$  on the basis of competent deduction from the premises  $p_1, \dots, p_n$ .
- (6) One knows each of  $p_1, \dots, p_n$ .
- (7) If one believes a conclusion on the basis of competent deduction from a set of premises one knows, one knows the conclusion (‘multi-premise closure’).
- (8) One knows something only if it has a high chance.

Treat (1)–(5) as an uncontentious description of the example. Relative to them, (6)–(8) form an inconsistent triad:

- (9) One knows  $p_1 \wedge \dots \wedge p_n$ . (from 5 to 7)
- (10) One does not know  $p_1 \wedge \dots \wedge p_n$ . (from 4, 8)

Which of (6)–(8) should we give up?

One might be tempted to reject (7) in light of this argument. The reason for this rejection might be the connection between knowledge and justification. Take any non-maximal (perhaps vague) threshold of justification and assume that in order to be justified in believing a proposition  $p$ , one’s total evidence  $e$  must sufficiently support  $p$ . Trivially, one might be (evidentially) justified in believing  $q$  since  $\Pr(q|e) > r$  (where  $r$  is the threshold of justification in the unit interval which falls short of 1) and might also be justified in believing  $q'$ , since it too surpasses the threshold:  $\Pr(q'|e) > r$ . And yet, as the case may be,  $\Pr(q' \wedge q|e) < r$ . Hence one will not be justified in believing what one competently deduces from one’s (evidentially) justified beliefs (by conjunction introduction) on one’s total evidence. If one thinks, that what holds for justification holds for knowledge, the natural reaction to the (1)–(10) contradiction, is to reject (7).<sup>10</sup>

But for Williamson there is a crucial difference between what holds for one’s justified beliefs and what holds for knowledge. On his account, one’s evidence is one’s

Footnote 9 continued

is open on defensible safety theories. For our purposes we will assume that safety does entail closure, or at least that the further conditions that Williamson imposes on knowledge guarantee its validity. For some interesting challenges the safety theory faces with regard to closure, see [Alspector-Kelly \(2011\)](#).

<sup>10</sup> We can add that the same holds for single-premise closure. Williamson’s account for justification in terms of evidence is, roughly, that one is justified (epistemically) in believing only those beliefs that one has evidence for (“An epistemically justified belief which falls short of knowledge must be epistemically justified by something; whatever justified it is evidence.” [Williamson 2000](#), p. 208). The “evidence for” relation as Williamson (as well as many others) understands it, requires that if  $e$  is evidence for  $p$ , then the probability of  $p$  given  $e$  is greater than the unconditional probability of  $p$  ([Williamson 2000](#), p. 187). Now, suppose one is evidentially justified in believing  $p$  (where one’s total evidence is  $e$ ). So  $\Pr(p|e)$  is very high but less than 1. Suppose further that there is some proposition  $q$  entailed by  $p$  the initial

total knowledge ( $E=K$ ), so the probability of anything that is known is 1 (since  $p$  is included in  $K$ —which is all that is known, for any known  $p$ ,  $\Pr(p|E) = \Pr(p|K) = \Pr(p|p \wedge K) = 1$ ). The principles governing justification, therefore, diverge significantly from the principles governing knowledge. No matter how many conjuncts one adds in the process of competent deduction, as long as the premises are known, the conclusion will have the same probability as the premises, namely, 1.

The epistemic probability that Williamson appeals to is by no means standard. But assuming for present purposes that there is no problem with the idea of objective prior probabilities as such, a consequence of Williamson's knowledge-evidence equation is that since the posterior probability of what is known is 1, the natural reaction to the puzzle—rejection of (7)—is not available. Even theorists who do not question *single premise* knowledge closure would be tempted to reject *multi premise* closure.<sup>11</sup> But, since adding known conjuncts by a deductive process of conjunction introduction will on Williamson's account always leave the probability of the conjunction unscathed ( $\Pr(q) = 1$ ), this natural line of reasoning is blocked for Williamson. For him, multi and single premise *knowledge* closure, stand or fall together.

So, rather than a rejection of (7), it is not surprising to find that Williamson rejects (8). This he achieves by drawing a distinction between objective chance (henceforth simply *chance*) and epistemic probability (henceforth *probability*), a distinction with quite far reaching consequences. The claim is that although the chance that the conjunction is true is low, its probability can be high, in fact in this case it is 1.

### 3 Lottery propositions

How does Williamson justify the sharp distinction between chance and probability? After all, it was more or less obvious to Lewis that there is a tight connection between

---

Footnote 10 continued

(non-evidential) probability of which is higher than  $\Pr(p|e)$ . It can be shown that for every proposition  $p$ , if  $\Pr(p|e) < 1$ , there is a proposition such as  $q$  that follows from  $p$  and has a lower probability given  $e$  than it has unconditionally  $\Pr(q|e) < \Pr(q)$  (for a formal proof of this point and further discussion see our (forthcoming)). Hence, although (by our assumption) one has evidence for  $p$ , one will not have evidence for a proposition that logically follows from it (assuming our/Williamson's principle of evidence). So, given Williamson's understanding of justification in terms of evidence (or knowledge), justification is not closed under entailment. Williamson's view, then, is already committed to a substantive separation of justification from knowledge and to the view that evidential justification is not closed. Notice that the proposition  $p$  cannot be taken as evidence on Williamson's account without a breach of the  $E=K$  principle. Interestingly, many theorists take justification closure to be just as intuitive as knowledge closure (see e.g. Gettier 1963). For them, Williamson's account will seem problematical. Moreover, non-closure of justification cannot be viewed as a local phenomena on his account, to rectify the situation one would need to reject major parts of Williamson's account since every non-maximally justified proposition entails propositions that follow from it that are not justified given his understanding of justification and evidence.

<sup>11</sup> The distinction is discussed in Hawthorne (2004, p. 33) with an insightful analysis of closure principles and the possibility of holding on to the single-premise version while discarding the multi-premise closure principle (e.g. 2004, pp. 141, 146, 154, 185–186).

the two. He famously claimed that objective chances should equal one's (admissibly) informed credence.<sup>12</sup>

Williamson's idea (which Hawthorne and Lasonen-Aarnio anticipate) is that what is objectively probable, i.e. chance, need not be represented in terms of close possible worlds. Suppose we represent a future event's chance of taking place as a branching out of a common history. When there is a chance that a different event, such as a quantum blip, will occur, no matter how slim the chance, there is a branch extending into the future from the actual world to one in which the blip takes place. Suppose we represent the conjunction  $q (= p_1 \wedge \dots \wedge p_n)$  in the above argument as a finite set of worlds with a common history up to a time  $t$ . We then have many branches extending from the set to worlds in which one of the events does not take place. Williamson's idea is that branching worlds are not necessarily close, where closeness is the central notion he employs to cash out his safety requirement (Williamson 2000, pp. 123–130). A safe belief is one that is true in all close possible world (Williamson 2009, pp. 324–325).<sup>13</sup> Since on his account  $S$  knows that  $q$  only if  $S$ 's true belief is safe from error, there can be no close worlds in which  $q$  does not hold no matter how slim the chance is of  $q$  being true. This is how radical a divorce Williamson is advocating between objective chance and epistemic probabilities.<sup>14</sup>

Yet this account of knowledge of what we can call quantum conjunctions (that is, conjunctions of very many propositions each having very high chance of being true adding up to low chance for the conjunction) runs into trouble. To see how, let us slightly modify a case presented by Hawthorne and Lasonen-Aarnio:

Consider extremely unlikely and bizarre 'quantum' events such as the event that a marble I drop tunnels through the whole house and lands on the ground underneath, leaving the matter it penetrates intact. (Hawthorne and Lasonen-Aarnio 2009, p. 94)

Let's imagine that we have a large yet finite amount of such marbles, say all currently existing marbles, such that on our best theories, it is quite likely (though not definite) that if all are dropped, at least one of them will in fact tunnel through the floor. As a matter of contingent fact, when we do drop them in some future time none of the marbles tunnel. Now the question we want to ask is this: Given that one has all and only information pertaining to the chances, does one know that none of the marbles will tunnel? In other words, given nothing but the story about likelihoods and assuming that it is in fact true, does one know that:

(11) For all existing marbles  $x$ , if  $x$  is dropped to the floor,  $x$  will not tunnel?

Whether he denies knowledge of (11) or allows such knowledge, Williamson's theory, it seems, is in trouble. Let us begin with the option of denying knowledge of (11). In

<sup>12</sup> Lewis (1986): "A Subjectivist's Guide to Objective Chance," and a modification of his view in "Humean Supervenience debugged" (1999).

<sup>13</sup> See also Williamson (2000, p. 102). Williamson stresses that closeness of worlds, cases, or belief episodes (there are several variants here) is not meant to be part of a reductive account of knowledge.

<sup>14</sup> To what extent can chance and probability diverge? Can the known chance be 0 while the epistemic probability is 1? As this question quickly gets entangled with issues that would take us far afield leave it as an open question for further deliberation.

Williamsonian terms, if (11) is not known this must be due to the fact that the belief in the truth of (11) is not adequately safe. This would mean that if we represent (11) as a long conjunction of propositions about dropped marbles tunneling, there is at least one conjunct that is not safely believed. But which? There is no forthcoming answer. It seems implausible that one would not know (11) on the basis of reasoning that every marble has extremely high chances of not tunneling (and assuming that all will in fact not tunnel).

Moreover, it is apparently false that there is some marble about which the belief that it will not tunnel is not safe.<sup>15</sup>

But there are further difficulties with denying knowledge of (11). If (11) is, under the circumstances we describe, unknown, then, it would seem, so must be (12):

(12) If *this* marble is dropped, it will not tunnel through the floor.

Assuming that we know a particular marble M will be dropped over a floor, does one know that the following is true?

(13) If M is dropped over *this* floor, M will not tunnel.

To avoid scepticism about future contingents Williamson must allow knowledge of (13). But if (11) and (12) are not known, it seems hard to explain how (13) could be.<sup>16</sup>

But even disregarding the behaviour of epistemic probabilities, it seems very strange to set the knowledge anywhere between (11) and (13)—either all are known, or none are.<sup>17</sup> Knowing none is scepticism, knowing all means knowing lottery propositions. Or so we will subsequently argue.

<sup>15</sup> If conjunctions of the form of (5) and (9) are known, it seems that (11) should be too, by the same sort of reasoning. Suppose we lay out the marbles and form a belief regarding each one (that it will not tunnel) and then add them up into a long conjunction from which (11) trivially follows.

<sup>16</sup> Let us note that there is a difference between cases of perception or memory and knowledge of future contingents and knowledge by induction. The former cases can be accounted for by an appeal to a form of epistemic externalism: if the evidence for there being a table in the room is seeing the table in the room, there is no mystery. When the fact itself is part of the evidence it is clear why the probability is 1. (Whether this form of externalism is plausible or defensible is another matter.) In the case of future contingents and induction it is clear that the evidence does not entail the known proposition. Supposing, as Williamson does, that one theory can account for these two sorts of knowledge is, perhaps, the heart of the matter. The epistemic situation seems to be very different when a proposition is entailed by the evidence (in which case assigning it epistemic probability 1 makes sense) and when it is not. There would be no need for induction if the inductive base entailed the known proposition. The idea that  $\text{Pr}(p|e) = 1$  when  $e$  does not entail  $p$ , to which Williamson is committed, seems to be the root of many of the problems. There is room then for bifurcating one's account of knowledge so that mathematical, perceptual, and knowledge based on memory is explained in Williamson's way, while inductive knowledge, for example, is explained by a more Lewisian conception of evidence. Some remarks by Williamson suggest that he may be more sympathetic to this kind of bifurcation than one might think:

“On the view above of evidence, when they constitute knowledge, they are part of our evidence. Moreover, they may constitute knowledge simply because perceiving counts as a way of knowing; that would fit the role of knowledge as evidence...I certainly did not perceive that your ticket did not win. There is no valid argument from the denial of knowledge in the lottery case to its denial in perceptual and other cases in which we ordinarily take ourselves to know” (Williamson 2000, p. 252).

<sup>17</sup> Of all the options the knowability of only (13), seems to be the worst. The relevant information about M being dropped should, if anything, lower its probability, since the information verifies its antecedent (making



We argued that preserving knowledge of claims about particulars while denying knowledge of related general claims is problematic on Williamson's view. A similar problem arises for the reverse situation, i.e. knowledge of general claims by induction from particular ones. I observe several ravens and note that they all appear black. Suppose that all ravens are black, and that at some point in the sequence of observations I come to know this (assuming induction is a method for obtaining knowledge). So at this point, I go from having a true belief that the probability of all ravens being black, such that  $\Pr(\text{For all } x \text{ (if } x \text{ is a raven, } x \text{ is black)}) < 1$ , to a state in which the probability equals 1. Although the transition from non-knowledge to knowledge is problematic on any account (Hume famously questions the very justification of induction), there is an added mystery in Williamson's account. My prior conditional probability of all ravens being black on my observing the next raven to be black was less than 1 and it increased steadily as evidence came in. But what is it about actually observing the next raven that changes the probabilities of all ravens being black to 1? Presumably, all theories of inductive knowledge will have to explain how before observing the raven I didn't know that all ravens are black, and now, after observing the relevant raven, I do. But for Williamson there is another difficulty stemming from the shift in probabilities. We are faced with the situation where we know that the proposition arrived at inductively does not follow from the evidence. The prior conditional probability of the hypothesis on the evidence is less than 1. Yet by receiving the evidence (which we know does not entail the proposition) we somehow arrive at probability 1 for that proposition. After observing the raven suddenly my evidence does entail the proposition. And this seems to get the order of explanation backwards. We want to know why the probability suddenly changes, i.e. why it goes from a conditional probability that is less than 1 to a conditional probability that is 1. The higher probability is supposed to tell us something about the knowledge, while here the knowledge explains the conditional probability change.<sup>18</sup>

Leaving aside the issue of induction, the idea that propositions like (13) are not known seems to be tangled with too many problems. It seems, then, that Williamson's theory would incline him to treat universal statements such as (11) in the same way

---

Footnote 17 continued

it more probable (and on Williamson's view, raising it to probability 1). Call the antecedent of (13)  $e$  and its consequent  $p$ . The probability of if  $e$  then  $p$ , is lower given  $e$ , than the unconditional probability of if  $e$  then  $p$ . On a standard Bayesian picture,  $\Pr(\neg(e \rightarrow p)|e) = \Pr((e \wedge \neg p) \wedge e)/\Pr(e) = \Pr(e \wedge \neg p)/\Pr(e) \geq \Pr(e \wedge \neg p) = \Pr(\neg(e \rightarrow p))$  and hence,  $\Pr((e \rightarrow p)|e) \leq \Pr(e \rightarrow p)$ . Given standard assumptions  $\Pr((e \rightarrow p)|e) < \Pr(e \rightarrow p)$ . Williamson rejects these assumptions after  $e$ , but given the case above this is to count against this rejection not against the standard assumptions. His rejection depends on counting  $\Pr((e \rightarrow p)|e)$ , as having probability 1 once  $e$  becomes known while the conditional probability before  $e$  becomes known is less than 1.

<sup>18</sup> To make things even clearer, the conditional probability before observing the raven includes the previously observed ravens up to the point before observing the knowledge-changing raven. After observing this raven the probability of the observation goes to 1 (as with all other standard accounts excluding those that use Jeffery Conditionalization). On Williamson's view the conditional probability of all ravens being black now includes (on the right side) not only the ravens observed so far but also the proposition that *all* ravens are black. The point in the main text is that including this proposition as part of the evidence is not explained by conditionalization which before observing the last raven was less than 1. Knowledge is appealed to in order to say why the conditional probability does not have the same probably before and after the observation.

he treats conjunctions of future contingents, namely as cases where, although chances are low, the epistemic probability is 1.

But suppose now that we have a lottery drawing in which 1 of a billion tickets will be drawn. Suppose further that all but one ticket have been sold and, coincidentally, it will be the one unsold ticket that will be the winner of the draw. So, for each of the sold tickets it is true both that its chances of losing are very high and that it will in fact lose. Is the belief that one of these tickets will lose safe? Williamson, like most epistemologists, thinks that lottery propositions are not known.<sup>19</sup> This is required for his explanation of the unassertability of lottery propositions in terms of their unknowability (Williamson 2000, pp. 224–229).<sup>20</sup> Merely having probabilistic reasons that a losing ticket in a large lottery will lose does not allow me to know, no matter how good my information is about the chances, that the ticket will not win. On Williamson's conception of safety, if belief in a conjunction is not safe, there must be at least one conjunct belief in which is not safe.<sup>21</sup> But as all conjuncts in this case are on a par, if belief in one is not safe, belief in any isn't. Thus, this would commit Williamson to a substantive distinction between quantum propositions—which are known, and lottery propositions—which are not. But what could be the difference? If I can't know that a lottery ticket is a loser, how can I know that a quantum blip will not occur, let alone know that the negation of a long disjunction about quantum events is true? If beliefs regarding falling marbles are safe, why not lottery beliefs?

To make the connection even tighter,<sup>22</sup> assume we match each of the lottery tickets to a marble dropping event (suppose we assign the numbers of the tickets to groups of marbles which are then dropped, and the winner is the holder of a ticket whose number matches a group of marbles of a tunnelling marble). It does not seem plausible in this case to say that although I know the marbles will not tunnel, I don't know my

<sup>19</sup> [H]owever unlikely one's ticket was to win the lottery, one did not know that it would not win, even if it did not ... No probability short of 1 turns true belief into knowledge (Williamson 2000, p. 117).

<sup>20</sup> On p. 255, Williamson (2000) connects the case of lotteries to the unassertability of the belief that one will not be knocked down by a bus tomorrow. It is hard to see how Williamson would separate this belief from beliefs about the non-occurrence of quantum events. If you don't know that you will not be knocked down by a bus, how can you know that a quantum blip will not happen in this instance? We stress here that the oddity of asserting lottery propositions is by no means the only reason Williamson does not accept knowledge that one's ticket will be a loser. Other reasons why Williamson cannot accept knowledge of lottery propositions include the very idea we are considering here—the idea that the evidential probability of any known proposition is 1. Surely one does not have evidence that one is going to lose if one has evidence that there is a slim chance that one will. Our claim here is that similarly one does not have conclusive evidence of the non-occurrence of a long conjunction of future events that have a low but positive chance of occurring. Another central idea that will be unattainable if lottery propositions are known, is the principle of multi-premise closure. Surly one cannot come to know that no sold lottery ticket will win and so one better buy the unsold ticket that is known (by closure) to be the winning ticket.

<sup>21</sup> Advocating what he calls the ordinary conception of safety, he claims:

Suppose that I am not safe from being shot. On the ordinary conception, it follows that there is someone *x* such that I am not safe from being shot by *x* (assume that if I am shot, I am shot by someone) (Williamson 2009, p. 327).

<sup>22</sup> The example in the main text preempts proposals such as that the distinction between lottery propositions and quantum proposition is that the latter but not the former are supported by nomological or scientific reasons. We thank an anonymous referee for this journal for pressing us on this point.

ticket is a loser. It is also implausible to claim that in such a set-up I know that no ticket will win the lottery. There are two possibilities here, either knowledge is lost by knowledge of the connection between the marbles and the lottery, or knowledge of the lottery propositions is gained by this known connection. If it is gained, then one can know that one's lottery ticket (as well as all the others) is a loser contrary to Williamson's conviction.<sup>23</sup> Loss of knowledge is equally dubious. Why would the fact that the quantum events are used as lottery mechanisms make them unknowable, if other quantum conjunctions are known?<sup>24</sup>

In general, it is hard to see why the world in which I win the lottery should be regarded more similar to those in which I lose, than the worlds in which a marble tunnels is to those in which none tunnel. The unclarity regarding the similarity (or closeness) relation at play in his account is related to a further lacuna in Williamson's presentation. Hardly ever does Williamson specify concrete instances of what by his lights would amount to knowledge. Can we know things by induction, or is the scenario in which they fail to be true too similar for such propositions to ever have probability 1? What can we know about the future? If he is not to slide too far on the way to skepticism, Williamson must allow that at least some knowledge of these sorts is possible. But then what could be the constraints on the similarity relations such that we get only the "good" cases and none of the "bad"?

A simple statement of our challenge is this: Are lottery propositions known or not? If they are, this would create problems for Williamson's thesis that knowledge is the norm of assertion (Williamson 2000, pp. 224–229) and commit Williamson to what is widely considered a very unfavourable position. Still worse, if lottery propositions are known multi-premise closure must be rejected (surely one does not know that no ticket holder has won before the lottery takes place). If, on the other hand, lottery propositions are not known, what is the relevant difference between them and quantum propositions? Specifically, if the lottery mechanism is just the quantum events, how can the latter be known while the former are not?

#### 4 Principal principle and practical deliberation

The divorce of epistemic probability from chance is intuitively problematic. Here is one way to give this intuition some substance. Since the chances can be known, in place of Williamson's (9), we might just as well have:

(9') One knows that the objective chance of  $p_1 \wedge \dots \wedge p_n$  is low.

<sup>23</sup> In private conversation, Timothy Williamson acknowledged that, unlike standard lottery situations, in such circumstances his theory is committed to knowledge of the lottery propositions.

<sup>24</sup> It is true that in the quantum lottery case there is no guaranteed winner, but we see no reason why this should make a difference to our argument. Particularly it seems hard to deny that one does not know that one's quantum lottery ticket is a loser if the chance it will win is greater than in an ordinary lottery. A derivative difference between the quantum lottery and the ordinary one is that with knowledge every losing ticket in the latter lottery there is greater probability that one of the other tickets will win, in the former lottery the quantum events are independent.

Is knowledge that the chance that some proposition is true is extremely low compatible with knowledge of this proposition? The answer to this question depends on the validity of a weakened version of Williamson's (8):

- (8') If one *knows* that the objective chance of a proposition  $q$  is low, one does not know  $q$ .

Williamson must reject (8'). Yet its rejection entails the truth of Moore-type future sentences of the form (for ease of exposition we use a higher order first person form):

- (14) I know that (the chance that  $q$  is true is low, but it will happen).

Given Williamson's knowledge account of assertion, the following instance of (14) is assertable: "the chance that my book contains no mistakes is very low, but it doesn't!" Strictly speaking, there is no contradiction in (14) or in any of its instances, just as there is no contradiction (at least no trivial contradiction) in any of the Moore-type sentences (many Moore sentences are true). With a further seemingly plausible principle, we can derive other variants of Moore sentences from (14) that sound even more odd.

- (15) If S knows that the objective chance of a future event is very low, then S knows that the future event might not take place.<sup>25</sup>

From (15) we can derive<sup>26</sup>:

- (16) S knows that ( $q$  might not be true but it is).

Note how odd that sounds in the case of future contingents: "I know it might not rain tomorrow, but I know it will." Obviously Williamson would prefer to regard (16) as unassertable. But since he is committed to the assertability of (14) this would mean he must take (15) to be false. The point of raising this issue (aside from the difficulties associated with rejecting (15) and (8')), is to shed more light on the radical rift Williamson is imposing between chances and epistemic probabilities. This is the focus of the present section.

Beyond intuition, however, there is a theoretical strain here. The sharp split between chance and epistemic probability conflicts—if not in letter, certainly in spirit—with the central idea motivating Lewis's Principal Principle which says:

Take some particular time – I'll call it "the present", but in fact it could be any time. Let C be a rational credence function for someone whose evidence is limited to the past and present – that is, for anyone who doesn't have any access to some very remarkable channels of information. Let P be the function that gives the present chances of all propositions. Let A be any proposition. Let E be any proposition that satisfies two conditions. First, it specifies the present chance of A, in accordance with the function P. Second, it contains no "inadmissible"

<sup>25</sup> It's hard to see how (16) could be false. Suppose I know that the chance for rain tomorrow is very low. Does it not seem adequate to say, at least, that I know that it might not rain tomorrow?

<sup>26</sup> Parentheses here are for avoiding scope ambiguity.

evidence about future history; that is, it does not give any information about how chance events in the present and future will turn out. ... Then the Principal Principle is the equation:  $C(A|E) = P(A)$  <sup>27</sup>

According to this principle one's credence in conjunctions with low chances should be just as low. Yet Williamson—along with other adherents of epistemic closure—is committed to the claim that one knows them. To know a proposition one must, presumably, believe it, which means that one must assign the proposition sufficiently high credence. It seems, then, that Williamson's desired conclusion requires abandoning the Principal Principle.

Williamson's response is to preserve the principle by allowing updating one's credences on evidence that Lewis regards "inadmissible". Specifically, by conditionalizing on the future contingents comprising the conjunction, which one is assumed to know, one's credence in the conjunction will be 1. Since for Lewis future contingents do not count as evidence such knowledge is inadmissible and is therefore not part of the formulation of the Principal Principle, which, as Williamson says, "is logically neutral as to the results of conditionalizing on inadmissible evidence, despite the forbidding connotations of the word "inadmissible".<sup>28</sup> For Williamson, all knowledge counts as evidence, including knowledge of the future. So one can conditionalize on this knowledge and update the credence assignments in accordance with the epistemic consequences of closure and the Principal Principle remains unviolated.

To the extent that Williamson's idea of conditionalizing on all knowledge succeeds,<sup>29</sup> it is a technical victory at best. Surely, even if Williamson manages to avoid violating the letter of Lewis's principle, he still undermines its spirit. The rationale behind the Principal Principle is that one should apportion one's credence in a proposition to what one (rationally) believes are the chances that proposition is true.

As I hope that the following questionnaire will show, we have some very firm and definite opinions concerning reasonable credence about chance. These opinions seem to me to afford the best grip we have on the concept of chance. Indeed, I am led to wonder whether anyone *but* a subjectivist is in a position to understand objective chance! (Lewis 1986, p. 84)

Thus, although Lewis's formulation of the Principal Principle is silent regarding inadmissible evidence, it is clear that its point is to articulate a tight connection between credence and chance. Williamson's position runs counter to this idea.

Let us make the point more explicit. In his reply to Hawthorne and Lasonen-Aarnio Williamson goes to great lengths to show that his view does not commit him to any implausible principle. But there is at least one highly plausible principle he seems to be forced to reject, call it the Weak Low Chance Principle:

<sup>27</sup> Lewis (1999, p. 238).

<sup>28</sup> Williamson (2009, p. 323).

<sup>29</sup> There are reasons to suspect that some form of "easy knowledge" is involved here, but we cannot develop this argument here. But see Sect. 3 and footnote 18.

WLC—If in  $w$ , at time  $t$ ,  $S$  knows that  $p$  has a low chance of being true,  $S$  does not know  $p$  at  $t$  in  $w$ .<sup>30</sup>

Given Williamson's divorce of epistemic probability from chance, he cannot endorse WLC. According to the safety theory of knowledge, one knows conjunctions of future contingents, for example, even when one knows their chances of being true are very slim.

To see just how problematic this is, consider the practical consequences of this commitment. Suppose the truth of some long conjunction of propositions is of crucial importance to you—if it is true you must  $\varphi$  and if false you must not. Now suppose you know each of the conjuncts is highly likely to be true and believe it to be true. Assuming that this belief is safely true, according to Williamson's theory you know the conjunction. But you also know that there is very high chance that the conjunction is false. Should you, or should you not  $\varphi$ ? Assuming knowledge is a rational basis for action, that it warrants taking something as a premise in practical deliberation (Hawthorne 2004, p. 29), it seems that Williamson's theory would entail that you ought to  $\varphi$ , despite the fact that you know there is very high chance that the conjunction is false. The reason for this, according to the theory is that although the chance that the conjunction is true is (and is known to be) low, since it is true and since you know this, you should act on this knowledge. But one can always turn the table on this argument. If you know that a proposition has high chance of being false you ought not act on it. After all, this too is knowledge. And acting on it is what the Principal Principle seems to suggest, again, if not by strict letter, than in spirit. Viewed from this perspective, a theory of knowledge that has consequences inconsistent with WLC is so much the worse for it, not the better. But even if it is discarded, we still face the problem of how one ought rationally to act when one has two incompatible directives stemming from one's knowledge. The point is not merely that we might not know what to take as a premise in practical deliberation because we don't know that we know. The point is that even if we were to have all relevant knowledge, we would still not know what to take as our premise.

## 5 Fallibilism

The problems we have raised for Williamson's safety theory of knowledge—the dilemma regarding lottery propositions and the relation between chance and credence—seem to both arise from the same feature of this theory, namely the attempt to treat human knowledge as, in a sense, infallible. While the precise formulation of epistemic fallibilism is a matter of contention, most epistemologists agree that human knowledge is, in some sense, fallible. Since the conception of knowledge as safely true belief has the consequence that what is known has an epistemic probability of 1, it is hard to see how any notion of fallibility can apply to knowledge under this conception (specifically if evidence is identified with knowledge). Consider some of the leading formulations.

<sup>30</sup> Compare this principle to Hawthorne's and Lasonen-Aarnio's Low Chance principle (2009, pp. 96–97) modified by Williamson (2009, p. 324).

“Fallibilism is the doctrine that someone can know that  $p$ , even though their evidence for  $p$  is logically consistent with the truth of not- $p$ ”.<sup>31</sup> For Williamson, remember, one’s knowledge is one’s evidence, so fallibilism in this formulation by Jason Stanley is not consistent with it. If you know  $p$ , then  $p$  is now part of your evidence and therefore, necessarily,  $p$  is true.<sup>32</sup> If one knows that  $p$ , one’s evidence is inconsistent with the falsity of  $p$ . Stewart Cohen’s formulation is even more directly at odds with Williamson’s theory: “a fallibilist theory allows that  $S$  can know  $q$  on the basis of  $r$  where  $r$  only makes  $q$  probable”.<sup>33</sup> Clearly, according to Williamson, this is false. Jim Pryor says: “a fallibilist is someone who believes that we can have knowledge on the basis of defeasible justification, justification that does not guarantee that our beliefs are correct.”<sup>34</sup> The notion of defeasibility employed here requires clarification. Nevertheless, it is clear that under any plausible theory, justification will be defeasible in the sense that future evidence can undermine it. Anything else would be an objectionable form of epistemic stubbornness.<sup>35</sup> The defeasibility associated with fallibilism must be something to do with the inconclusive nature of that on which ordinary knowledge is based and maintained. This, presumably, is what Pryor means when he speaks of “justification that does not guarantee that our beliefs are correct.” Richard Feldman articulated this idea more explicitly. Fallibilism, he says, is the view that “it is possible for  $S$  to know that  $p$  even if  $S$  does not have logically conclusive evidence to justify believing that  $p$ ”.<sup>36</sup> As he explains, this amounts to the claim that knowledge can be had based on less than deductive inference, that one’s evidence need not entail what is known. But if knowledge is safely true belief, belief that has epistemic probability 1 on one’s evidence, then it *is* guaranteed—evidence *is* conclusive and entails what is known (at least epistemically).

There is, perhaps one kind of fallibilism we can think of that is compatible with Williamson’s safety theory of knowledge.<sup>37</sup> Since epistemic probabilities are for Williamson divorced from objective chances, it is consistent with his theory that one can know things which have less than perfect chance of being true. Indeed, as we have seen, this is a central desiderata of his theory. It may be claimed therefore that our epistemic fallibilism consists in the fact that we can know propositions which—objectively speaking—have a high chance of being false. The important thing to notice about this proposal, however, is that it allays the epistemic bite of fallibilism. As the attempts to define fallibilism all indicate, the idea that knowledge is fallible is supposed to capture something about the relation between knowledge and the evidence or

<sup>31</sup> Stanley (2005, p. 127).

<sup>32</sup> Perhaps Stanley [who is sympathetic to the idea that knowledge has probability 1—see Hawthorne and Stanley (2008)] could change the formulation thus: Fallibilism is the doctrine that someone can know that  $p$ , even though the evidence one had for  $p$  and by which one came to know that  $p$ , is logically consistent with the truth of not- $p$ .

<sup>33</sup> Cohen (1988, p. 91).

<sup>34</sup> Pryor (2000, p. 518).

<sup>35</sup> This kind of epistemic stubbornness is the one that concerns Saul Kripke in his famous puzzle, in “On Two Paradoxes of Knowledge” (Kripke 2011).

<sup>36</sup> Feldman (1981, p. 266).

<sup>37</sup> We thank Krista Lawlor for drawing our attention to this possibility.

justification on which it is (or can be) based. This is lost in the definition we have proposed on Williamson's behalf. With this definition all that the fallibility of knowledge comes to is the uninteresting claim that we can know truths that have some chance of being false, although they aren't, and our evidence guarantees that they aren't. If this is as robust a notion of fallibilism as Williamson can endorse, a strong and prevalent fallibilist intuition is left unsatisfied and the problems discussed in previous sections remain intact. The low chance of the proposition's truth suggests a straightforward way in which what is known can fail to be true. But, for Williamson, this possibility of error, this chance of falsity, carries no epistemic weight (otherwise the probability of the conjunction will not be 1 false and closure will fail). But it is precisely epistemic fallibility that is intuitively associated with knowledge. At the same time, endorsing infallibilism while acknowledging the objective high chance of error is also unsatisfying. Picking up an argument recently proposed by Martin Smith, it is hard for Williamson's theory to explain why the belief that a screen displays a blue background is regarded as evidentially fallible when held by someone who knows that it displays blue 999,999 times out of a million and infallible when held by someone who observed the blue screen, but who has made more than one error for every million visual observations. As Smith argues, "the only reasons that the evidential probability of P is taken to be higher for [the latter] than for [the former] is because [the latter's] belief is taken to qualify as knowledge whereas [the former's] is not."<sup>38</sup> Williamson, in other words, is putting the carriage before the horses—a belief is evidentially infallible because it is known, instead of being known in virtue of being infallible. The problem we have seen with induction (Sect. 3) seems to resurface here.

In addition, Williamson's infallibilism has particularly unsettling consequences with respect to belief revision. Surely there are psychological facts about a subject that might undermine knowledge of  $p$  such as ceasing to believe it. But loss of knowledge and change of belief is also sometimes rationally required, specifically, when proper counter-evidence presents itself. On Williamson's account, once something is known such change is never rationally mandated. This is because one can always take one's knowledge that  $p$  as evidence ruling out any evidence to the contrary. To illustrate this point consider Williamson's own example:

I put exactly one red ball and one black ball into an empty bag, and will make draws with replacement. Let  $h$  be the proposition that I put a black ball into the bag, and  $e$  the proposition that the first ten thousand draws are all red. I know  $h$  by standard combination of perception and memory, because I saw that the ball was black and I put it into the bag a moment ago. Nevertheless, if after ten thousand draws I learn  $e$ , I shall have ceased to know  $h$ , because the evidence which I shall then have will make it too likely that I was somehow confused about the colours of the balls. (Williamson 2000, p. 205)

This surely seems to be the rational way to go. But  $e$  is never simply given. After ten thousand draws one is faced with two conflicting pieces of information,  $h$  and  $e$ . If  $h$  is known and therefore has probability 1, it would be just as reasonable to question

<sup>38</sup> Smith (2010, p. 18).



one's memory, which, presumably, is the basis of one's belief in  $e$ . One can always conditionalize on one's knowledge (evidence) that there is a black ball in the bag and conclude that one is confused about  $e$  ( $\Pr(e|h) = 1/2^{10,000}$ ), not about  $h$ . This is a further sense in which Knowledge seems to be fallible in a more substantial way than Williamson's view allows. For all intents and epistemological purposes, then, Williamson's theory of knowledge entails a form of infallibilism.

## 6 Conclusion

We have shown that ordinary, fallibilist theories of knowledge cannot maintain the epistemic role of inference, even of a very basic kind. Against this background, Williamson's safety theory of knowledge was presented as having the significant virtue of preserving the idea that knowledge is expanded by deductive inferences such as modus ponens. Nonetheless, we have argued, this gain comes at great cost. To be sure, we have said nothing to conclusively refute Williamson's theory of knowledge, nor to exclude the possibility that other ways might be devised to overcome the problem of uncertainty accumulation. Choosing between competing theories in such matters is more often a matter of balancing costs against benefits than of conclusive refutation and proof. Our contribution therefore is in articulating some of the costs incurred by what seems to be the most promising proposal on offer. We have shown that this theory faces great difficulty in accounting for lottery propositions, whether it takes them to be knowable or not. We also showed that it must give up a highly plausible idea underlying Lewis's widely endorsed Principal Principle—the idea that one should apportion one's belief in a proposition to what one believes (or knows) about the chances of it being true. Problems regarding rational decision surface when knowledge of chances is divorced from other forms of knowledge regarding the very same events, we have argued. All this in addition to and following from the fact that accepting Williamson's theory entails the rejection of a central intuition endorsed by a majority of epistemologists, namely that knowledge is fallible.

**Acknowledgments** For helpful comments many thanks to the participants of the Van Leer Jerusalem Institute's Epistemology Study Group and Stockholm University's Graduate Seminar in Theoretical Philosophy. For very helpful comments and suggestions many thanks to Stewart Cohen, John Hawthorne, Karl Karlander, Krista Lawlor, Avishai Margalit, Jim Pryor, Oron Shagrir, Jonathan Yaari, Jonas Åkerman and two anonymous referees for this journal. Special thanks to Peter Pagin and Timothy Williamson for very helpful discussions.

## References

- Alspector-Kelly, M. (2011). Why safety doesn't save closure. *Synthese*, 183(2), 127–142.
- Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives*, 2, 91–123.
- Christensen, D. (2004). *Putting logic in its place: Formal constraints on rational belief*. New York: OUP.
- Feldman, R. (1981). Fallibilism and knowing that one knows. *Philosophical Review*, 90(2), 266–282.
- Fumerton, R. (2006). *Epistemology*. Malden, MA: Blackwell.
- Gettier, E. L. (1963). Is justified true belief knowledge?. *Analysis*, 23, 121–123.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: OUP.
- Hawthorne, J., & Lasonen-Aarnio, M. (2009). Knowledge and objective chance. In P. Greenough & D. Pritchard (Eds.), *Williamson on knowledge* (pp. 92–108). Oxford: Oxford University Press.

- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105(10), 571–590.
- Kripke, S. (2011). On two paradoxes of knowledge. In *Philosophical troubles, collected papers* (Vol. I). Oxford: OUP.
- Kyburg, H. (1961). *Probability and the logic of rational belief*. Middletown, CT: Wesleyan University Press.
- Kyburg, H. (1970). Conjunctivitis. In M. Swain (Ed.), *Induction, acceptance and rational belief* (pp. 55–82). Dordrecht: Reidel.
- Lewis, D. (1986). A subjectivist's guide to objective chance. In *Philosophical papers* (Vol. II, pp. 83–113). Oxford: Oxford University Press.
- Lewis, D. (1999). Humean supervenience debugged. (Reprinted from *Papers in metaphysics and epistemology*, by D. Lewis, Ed., Cambridge: Cambridge University Press).
- Makinson, D. (1965). The paradox of the preface. *Analysis*, 25, 205–207.
- Pagin, P. (1990). Review of Roy Sorenson, *Blindspots*, Clarendon Press, Oxford 1988. *History and Philosophy of Logic*, 11, 243–245.
- Pryor, J. (2000). The skeptic and the dogmatist. *Noûs*, 34(4), 517–549.
- Sharon, A., & Spectre, L. (forthcoming). Evidence and the openness of knowledge. *Philosophical Studies*.
- Smith, M. (2010). What else justification could be. *Noûs*, 44(1), 10–31.
- Stanley, J. (2005). Fallibilism and concessive knowledge attributions. *Analysis*, 65, 126–131.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: OUP.
- Williamson, T. (2009). Reply to John Hawthorne and Maria Lasonen-Aarnio. In P. Greenough & D. Pritchard (Eds.), *Williamson on knowledge*. Oxford: Oxford University Press.