

## Improbable Knowing\*

(to appear in T. Dougherty, ed., *Evidentialism and its Discontents*, Oxford: Oxford

University Press)

Timothy Williamson

University of Oxford

How much can we turn the screw on counter-examples to the KK principle? The principle, also sometimes called “positive introspection”, says that if one knows that P, one knows that one knows that P. It is widely, although not universally, acknowledged that the KK principle is false, and not just for the boring reason that one can know that P without having formed the belief that one knows that P. One can know that P, and believe that one knows that P, without knowing that one knows that P, because one is not in a strong enough epistemic position to know that one knows that P (Williamson 2000: 114-30). One’s epistemic position is weaker with respect to the fact that one knows that P than it is with respect to the fact that P itself. But how *much* weaker can one’s epistemic position be with respect to the fact that one knows that P than it is with respect to the fact that P?

Of course, if one does know that P, one’s epistemic position with respect to the fact that one knows that P cannot be so bad that one positively knows that one does not know that P, as opposed to merely failing to know that one knows that P, for knowing is

factive: one can know only what is indeed the case. But we can ask a more graded question in terms of evidential probabilities. Given that one knows that P, how low can the probability, on one's own present evidence, that one knows that P sink? It will be argued here that it can sink arbitrarily close to 0. At the limit, the probability on one's evidence that P can be 1 while the probability on one's evidence that one knows that P is 0. The difference between the probabilities can be as large as probabilistic differences can go.

One way of arguing for such a conclusion follows a traditional way of arguing against the KK principle, by using familiar examples that support some form of externalism about knowledge (Lemmon 1967). For instance, the unconfident examinee answers questions on English history under the impression that he is merely guessing. In fact, his answers are correct, and result from lessons on it that he has completely forgotten he ever had (Radford 1966). The example can be so filled in that it is extremely improbable on the examinee's evidence that he had any such lessons, or any other access to the relevant knowledge of English history; nevertheless, he does know the historical facts in question. That description of the example may well be right. Dialectically, however, it has the disadvantage that those of a more internalist bent may simply deny that the examinee knows. This paper develops a more systematic, structural way of arguing for the realistic possibility of knowing when it is extremely improbable on one's evidence that one knows, a way that does not depend on specifically externalist judgments about cases.

On the resulting view, one's evidence can be radically misleading about one's own present epistemic position. If the rationality of an action depends on one's epistemic

position, then one's evidence can be radically misleading about the rationality of the various actions available to one. Such phenomena will be used to cast light on some epistemological puzzles.

1. It is useful to explore the issues within a framework taken from the standard possible worlds semantics for epistemic logic, introduced by Hintikka (1962). Such a formal framework keeps us honest, by making it straightforward to check whether our descriptions of examples are consistent and what their consequences are, and by facilitating the identification of structurally appropriate models. Of course, we must also consider whether the mathematical models we use are realistic on their intended epistemic interpretation in the relevant respects. It will be argued below that the respects in which they are idealized are consistent with the uses to which the models are here being put.

We recall some basic features of possible worlds models for epistemic logic. For present purposes, we can make two convenient simplifications. First, we need only consider one agent at a time. Although important analogues of the failure of the KK principle arise for the interpersonal case too (Williamson 2000: 131-4), we can ignore such complications here. Second, we can ignore the strictly semantic aspect of possible world models, by discussing propositions rather than sentences; the resultant structures are *frames*.

Given those simplifications, a frame is just an ordered pair  $\langle W, R \rangle$ , where  $W$  is a set and  $R$  a binary relation on  $W$  (a set of ordered pairs of members of  $W$ ). Informally, we think of  $W$  as the set of relevant worlds, and  $R$  as a relation of epistemic possibility

between them, where a world  $x$  is epistemically possible in a world  $w$  ( $Rwx$ ) if and only if, for all one knows in  $w$ , one is in  $x$ , in other words, whatever one knows in  $w$  is true in  $x$  (where ‘one’ refers to the relevant agent and the present tense to the relevant time).

Correspondingly, we think of the subsets of  $W$  as propositions; a proposition  $p \subseteq W$  is true in a world  $w$  if and only if  $w \in p$ . Obviously the conjunction of two propositions is their set-theoretic intersection, the negation of a proposition is its complement in  $W$ , and so on.

For  $p \subseteq W$ , we define:

$$Kp = \{w \in W: \forall x \in W (Rwx \rightarrow x \in p)\}$$

Informally,  $Kp$  is to be the proposition that one knows  $p$ .

On its intended reading, the definition of  $K$  presupposes that one knows something if and only if it is true in all epistemic possibilities for one, that is, in all worlds in which whatever one knows is true. This involves assumptions about the agent’s logical omniscience, in two ways. First, any treatment of  $K$  as a function from propositions to propositions automatically requires that if  $p$  is the same proposition as  $q$  then  $Kp$  is the same proposition as  $Kq$ . Since propositions are being treated as sets of worlds, this means that if  $p$  is true in the same worlds as  $q$ , then  $Kp$  is true in the same worlds as  $Kq$ . Moreover, since truth in a world respects the usual truth-functions — a conjunction is true in a world if and only if all its conjuncts are true in that world, the negation of a proposition is true in a world if and only if the original proposition is not true in that world, and so on — truth-functionally equivalent propositions are identical, so both are known or neither is; in particular, every truth-functional tautology is known if any is. That first form of logical omniscience is independent of the specifics of the right-hand side of the definition. The second form depends on those specifics, but unlike the


first applies whenever premises  $p_1, \dots, p_n$  entail a conclusion  $q$  ( $p_1 \cap \dots \cap p_n \subseteq q$ ), even if the conclusion is not identical with any of the premises; it says that, in those circumstances, knowing the premises entails knowing the conclusion ( $Kp_1 \cap \dots \cap Kp_n \subseteq Kq$ ), for if each premise is true in all worlds in which whatever one knows is true, the conclusion is also true in all worlds in which whatever one knows is true.

On most views of propositional attitudes, logical omniscience is an extreme idealization. Someone can know that  $P \rightarrow P$  without knowing that  $((P \rightarrow Q) \rightarrow P) \rightarrow P$ , even though both formulas are tautologies. For present purposes it is a harmless one, for if the total evidence of a logically omniscient agent can be radically misleading about their epistemic position, our own abject failure of logical omniscience will not save us from the same fate. We will return to this issue towards the end of the paper.

One effect of full logical omniscience is that for each world  $w$  there is the strongest proposition known by the agent in  $w$ ,  $R(w)$ , in the sense that it is known in  $w$  and entails every proposition that is known in  $w$ . We can define  $R(w) = \{x \in W: Rwx\}$ . Then a proposition  $p$  is known in  $w$  if and only if  $p$  follows from  $R(w)$ ; more formally,  $w \in Kp$  if and only if  $R(w) \subseteq p$ , by definition of  $K$ .

One constraint on the epistemic possibility relation will be maintained throughout.  $R$  is reflexive; every world is epistemically possible in itself, because knowledge entails truth: whatever one knows in a world is true in that world. Consequently,  $w \in R(w)$  for any world  $w$ , and  $Kp \subseteq p$  for any proposition  $p$ .

Now recall the formal structure of counterexamples to the KK principle in such frames for epistemic logic. Here is a toy example.  $W$  is a three-member set  $\{x, y, z\}$ :



In the diagram,  $R$  holds between worlds just in case they are identical or neighbours; thus  $R$  is both reflexive and symmetric.  $R$  is not transitive, because  $Rxy$  and  $Ryz$  but not  $Rxz$ .

The strongest things known in each world are these:  $R(x) = \{x, y\}$ ;  $R(y) = \{x, y, z\}$ ;  $R(z) = \{y, z\}$ . In effect, if one is at one of the endpoints, what one knows is that one is not at the other endpoint; if one is at the midpoint, one knows nothing non-trivial about one's position. Now let  $p = \{x, y\}$ . Then  $Kp = \{x\}$ : one knows  $p$  in  $x$  because  $p$  is true in all worlds epistemically possible in  $x$ ; one does not know  $p$  in  $y$  because  $p$  is false in  $z$ , which is epistemically possible in  $y$ . Consequently  $KKp = K\{x\} = \{\}$ : one does not know  $Kp$  in  $x$  because  $Kp$  is false in  $y$ , which is epistemically possible in  $x$ . Thus the KK principle fails in  $x$ , because  $Kp$  is true and  $KKp$  false there.

As is well known, the non-transitivity of  $R$  is necessary and sufficient for a frame to contain a counterexample to the KK principle. For consider any frame  $\langle W, R \rangle$ . Suppose that  $R$  is non-transitive. Thus for some  $x, y, z$  in  $W$ ,  $Rxy$  and  $Ryz$  but not  $Rxz$ . By definition,  $KR(x)$  is true in  $x$ .  $KR(x)$  is not true in  $y$ , because  $Ryz$  and  $R(x)$  is not true in  $z$  (since not  $Rxz$ ). Therefore  $KKR(x)$  is not true in  $x$ , because  $Rxy$ . Thus the KK principle fails in  $x$ . Conversely, suppose that there is a counter-example to the KK principle in  $\langle W, R \rangle$ , say in  $x \in W$ . Thus for some  $p \subseteq W$ ,  $Kp$  is true in  $x$  and  $KKp$  false in  $x$ . By the latter, for some  $y \in W$ ,  $Rxy$  and  $Kp$  is false in  $y$ , so for some  $z \in W$ ,  $Ryz$  and  $p$  is false in  $z$ . But not  $Rxz$ , otherwise  $Kp$  is false in  $x$ , contrary to hypothesis. Thus  $R$  is non-transitive.

Of course, the existence of such non-transitive frames for epistemic logic does not by itself establish that there are counterexamples to the KK principle on its intended interpretation, for it remains to be shown that these mathematical structures represent genuinely possible epistemic situations. Before we turn to such matters, however, we must first enhance the frames with probabilistic structure, so that we can model issues about the probability on one's evidence that one knows something.

2. In adding probabilities to a frame  $\langle W, R \rangle$ , the account of evidential probability in Williamson (2000: 209-37) will be followed. We start with a prior distribution  $\text{Prob}_{\text{prior}}$  over propositions. Thus we can take a probabilistic epistemic frame to be an ordered triple  $\langle W, R, \text{Prob}_{\text{prior}} \rangle$ , where  $W$  and  $R$  are as before and  $\text{Prob}_{\text{prior}}$  is a probability distribution defined over subsets of  $W$ .

In the frames considered in detail below,  $\text{Prob}_{\text{prior}}$  always takes the particularly simple form of a uniform distribution over the subsets of a finite set  $W$ , in the sense that every world has the same probabilistic weight as every other world. Thus, where  $|p|$  is the cardinality of  $p \subseteq W$ ,  $\text{Prob}_{\text{prior}}(p) = |p|/|W|$ . It is not suggested that non-uniform or infinite probability distributions are in any way illegitimate. However, if a non-uniform distribution were used to illustrate the epistemic phenomena in question, they might look like artefacts of gerrymandering. Similarly, if  $W$  were infinite, the phenomena might look like paradoxes of infinity, given the complications of probability distributions over infinite sets. It is therefore best to use a uniform prior distribution over a finite space where possible.

For such uniform prior distributions, every nonempty subset of  $W$  has nonzero probability. We can therefore unproblematically define prior conditional probabilities by ratios in the usual way:  $\text{Prob}_{\text{prior}}(p \mid q) = \text{Prob}_{\text{prior}}(p \cap q) / \text{Prob}_{\text{prior}}(q)$  whenever  $q$  is nonempty and so  $\text{Prob}_{\text{prior}}(q) > 0$ .

The evidential probability of a proposition in a world  $w$  is identified with its probability conditional on one's total evidence in  $w$ . One's total evidence in  $w$  can in turn be identified with the total content of what one knows in  $w$  (Williamson 2000: 184-208). In a frame  $\langle W, R \rangle$ , the total content of what one knows in  $w$  is just  $R(w)$ . Since  $w \in R(w)$ ,  $R(w)$  is always nonempty, so probabilities conditional on  $R(w)$  are always well-defined. So if  $\text{Prob}_w(p)$  is the evidential probability in  $w$  of a proposition  $p$ :

$$\text{Prob}_w(p) = \text{Prob}_{\text{prior}}(p \mid R(w)) = \text{Prob}_{\text{prior}}(p \cap R(w)) / \text{Prob}_{\text{prior}}(R(w))$$

Thus in finite uniform frames, the evidential probability in  $w$  of  $p$  is simply the proportion of epistemically possible worlds in  $w$  in which  $p$  is true.

We can locate propositions about evidential probabilities in the frame. For instance, the proposition  $[Pr(p) = c]$  that the evidential probability of  $p$  is the real number  $c$  is simply  $\{w \in W : \text{Prob}_w(p) = c\}$ , and similarly for inequalities involving evidential probabilities. Thus propositions about evidential probabilities will themselves have evidential probabilities.

Let  $\langle W, R \rangle$  be the three-world toy example from the previous section. As before,  $p$  is  $\{x, y\}$ , so  $Kp$  is  $\{x\}$ . Let  $\text{Prob}_{\text{prior}}$  be the uniform distribution for  $W$ , so  $\text{Prob}_{\text{prior}}(\{x\}) = \text{Prob}_{\text{prior}}(\{y\}) = \text{Prob}_{\text{prior}}(\{z\}) = 1/3$ . Hence  $\text{Pr}_x(Kp) = 1/2$ , since  $Kp$  is true in just one of the




two worlds that are epistemically possible in  $x$ . Thus in  $x$ , even though one knows  $p$ , the probability on one's evidence that one knows  $p$  is no more than 50-50. To say that the probability on one's evidence that one knows  $p$  is just  $\frac{1}{2}$  is already to say something much worse about the epistemic status for one of the proposition that one knows  $p$  than merely to say that one does not know that one knows  $p$ .

Let  $\langle W, R, \text{Prob}_{\text{prior}} \rangle$  be any probabilistic epistemic frame where  $W$  is finite,  $R$  is reflexive and  $\text{Prob}_{\text{prior}}$  is uniform. Suppose that the KK principle fails in the frame. Thus for some  $w \in W$  and  $p \subseteq W$ ,  $Kp$  is true in  $w$  while  $KKp$  is false in  $w$ . Thus for some  $x \in R(w)$ ,  $Kp$  is false in  $x$ . Since  $\text{Prob}_{\text{prior}}(\{x\}) > 0$ ,  $\text{Prob}_w(Kp) < 1$ . Thus wherever the KK principle fails in such models, one knows something although the probability on one's evidence that one knows it is less than 1. By contrast, if the KK principle holds in a frame, if  $Kp$  is true in  $w$  then  $KKp$  is true in  $w$ , so  $R(w) \subseteq Kp$ , so  $\text{Prob}_w(Kp) = 1$ : whenever one knows something, the probability on one's evidence that one knows it is 1. Indeed, in such models knowing  $p$  is equivalent to its having probability 1 on one's evidence;  $Kp = [Pr(p)=1]$ . Thus the KK principle is equivalent to the principle that if the evidential probability of  $p$  is 1, then the evidential probability that the evidential probability of  $p$  is 1 is itself 1.

In frames where  $W$  is finite,  $R$  is reflexive and  $\text{Prob}_{\text{prior}}$  is uniform, how low can  $\text{Prob}_w(Kp)$  be when  $Kp$  is true in  $w$ ? When one knows something, how low can the probability that one knows it be on one's evidence? The probability can be any rational number whatsoever strictly between 0 and 1, so it can be arbitrarily close to 0. To see this, let  $m/n$  be any rational number such that  $0 < m/n < 1$ , where  $m$  and  $n$  are positive integers, so  $0 < m < n$ . We construct a model of the required kind with some worlds in

which  $Kp$  is true while the evidential probability of  $Kp$  is  $m/n$ . The idea is simple: take the three-world toy model used to illustrate the failure of the KK principle in section 1, multiply the world  $x$  by  $m$ , yielding  $m$  mutually indiscernible copies, and the world  $y$  by  $n-m$ , yielding  $n-m$  mutually indiscernible copies. Thus  $W = \{x_1, \dots, x_m, y_{m+1}, \dots, y_n, z\}$ .  $R$  is reflexive and symmetric; for any  $i, i^*, j, j^*$  where  $1 \leq i, i^* \leq m < j, j^* \leq n$ :  $Rx_i x_{i^*}$ ,  $Rx_i y_j$ ,  $Ry_j y_{j^*}$ , and  $Ry_j z$ , but not  $Rx_i z$ ; diagrammatically:



If  $p = \{x_1, \dots, x_m, y_{m+1}, \dots, y_n\}$ , for  $1 \leq i \leq m$ ,  $R(x_i) = p$ . Thus  $Kp = \{x_1, \dots, x_m\}$ .

Consequently, while in  $x_i$   $Kp$  is true, it is true in only  $m$  of the  $n$  epistemically possible worlds. Since  $\text{Prob}_{\text{prior}}$  is uniform,  $\text{Prob}_{x_i}(Kp) = m/n$ , as required. In particular, for  $m = 1$ ,  $\text{Prob}_{x_i}(Kp) = 1/n$ , which goes to 0 as  $n$  goes to infinity. By using non-uniform prior probability distributions or infinite sets of worlds we could construct similar models in which  $Kp$  is true while actually having evidential probability 0, but such refinements are unnecessary here.<sup>1</sup>

In the model just illustrated,  $m/n$  is the highest evidential probability that  $Kp$  attains anywhere, since  $\text{Prob}_{y_j}(Kp) = m/(n+1)$  and  $\text{Prob}_z(Kp) = 0$ . Thus the proposition  $[Pr(Kp) \leq m/n]$  is true at every world in the model. Consequently, so is the proposition

$K^k[Pr(Kp) \leq m/n]$ , where  $k$  is a natural number and  $K^k$  means  $k$  iterations of  $K$  (thus  $K^0 q$  is  $q$  and  $K^{k+1} q$  is  $K^k K q$ ). In other words, knowing  $p$  is compatible not just with the probability on one's evidence that one knows  $p$  being close to 0, but even with knowing, and knowing that one knows, and knowing that one knows that one knows, ... that the probability on one's evidence that one knows  $p$  is close to 0.

We have added evidential probabilities to epistemic models in a way that embodies several strong assumptions. In particular, one's total evidence was equated with the total content of one's knowledge, and probabilities on that evidence were calculated by conditionalizing a prior probability distribution on it. These assumptions are defensible (Williamson 2000), but can of course be challenged. However, someone who denied that they always hold would not be thereby committed to rejecting their present applications. For they are being used to argue that a specific phenomenon *can* occur, not that it *always* occurs. That requires that the relevant models *can* be instantiated by genuine epistemic phenomena, not that all genuine epistemic phenomena have the structure of those models. Indeed, the assumptions at issue should make it *harder* to construct models with the target phenomenon, which involves a sort of tension between knowing and the evidential probability of knowing. For what is most distinctive about the present approach is the intimate connection it postulates between evidential probabilities and knowledge; the assumptions cramp attempts to arrange the tension between them. By contrast, an approach that allowed more independence between evidential probabilities and knowledge would have much more scope to arrange the tension, by varying the evidential dimension independently of the knowledge dimension or *vice versa*. Similarly, allowing non-uniform prior probability distributions or infinite sets of worlds would give

far more scope for arranging odd probabilistic phenomena, for example by giving special weight to ‘bad’ worlds. If the target phenomenon occurs even under the unhelpful restrictive conditions postulated by the present approach to evidential probabilities, it is robust. A less restrictive approach could hardly rule out the models already constructed. The challenge to opponents is to motivate an approach that is *more* restrictive in some relevant way.

So far, however, we have been working at the level of formal models, without any positive argument that they represent genuinely possible epistemic situations. We now turn to that task, and provide a much more realistic, only slightly schematized description of a mundane type of epistemic situation that exemplifies the target phenomenon.

3. Imagine a plain, unmarked circular dial with a single pointer, like the hand of a modernist clock-face designed with an eye to the appearance not the reality of functional efficiency. The stationary hand can point at any one of  $n$  equally spaced unmarked positions on the perimeter of the dial; it is positioned at random. We measure distances between positions by the minimum number of steps needed to go from one to the other (clockwise or anti-clockwise). Number the positions  $0, \dots, n-1$  clockwise from the top. For simplicity, we individuate ‘worlds’ (mutually exclusive and jointly exhaustive relevant circumstances) just by the position of the hand. Thus there are exactly  $n$  worlds  $w_0, \dots, w_{n-1}$ , where in  $w_i$  the hand points at position  $i$ . We measure distances between worlds by the corresponding distances between positions. An *interval* of positions corresponds to a time interval; it is a nonempty proper subset of the set of positions such that the hand goes through every position in the set without going through any position

not in the set (intervals are ‘connected’). An *endpoint* of an interval is a member next to a non-member. Any interval with at least two members has two endpoints. An interval with an odd number of members has a unique *midpoint*, equidistant from its endpoints. Given the natural one-one correspondence between worlds and positions, the terms ‘interval’, ‘endpoint’ and ‘midpoint’ can be applied just as well to sets of worlds as to sets of positions.

Now imagine that you are looking at the dial from a fixed point of view equidistant from all points on the perimeter, and that you have no independent information about the hand’s position. You can make some discriminations between positions, and correspondingly between worlds, but the difference between neighbouring positions is well below your threshold of discrimination. We may assume that your capacity to discriminate between positions depends only on their relative distance; thus if world  $w$  is at least as close to world  $x$  as world  $y$  is to world  $z$ , then you can discriminate  $w$  from  $x$  only if you can also discriminate  $y$  from  $z$ . Consequently, if you are in fact in world  $w$ , the worlds that for all you know you are in (the epistemically possible worlds) are those at most  $h$  steps from  $w$ , for some natural number  $h$ ;  $h$  is greater than 0, otherwise your discrimination would be perfect. We can regard  $h$  as the width of the margin for error you require in order to know something in the model (Williamson 2000: 125-34); it is a constant with respect to the given model. More formally, let  $R$  be the epistemic accessibility relation; then for all worlds  $w$  and  $x$ ,  $Rwx$  if and only if the distance between  $w$  and  $x$  is at most  $h$ . Thus  $R$  has both reflective and rotational symmetry. For any world  $w$ ,  $R(w)$  (the set of worlds epistemically possible in  $w$ ) is not the whole of  $w$ , otherwise you could not make any discriminations at all. Thus  $R(w)$  is an

interval with  $w$  as its midpoint and  $2h+1$  members. For instance, if we imagine the distance  $h$  around the circumference as five minutes on a clock,  $R(w)$  corresponds to the period starting five minutes before the time  $w$  and ending five minutes after  $w$ .

We can prove that  $R(w)$  is known *only* at  $w$ . For suppose that  $R(w)$  is known at a world  $x$ . Since  $R(x)$  is the strongest proposition known at  $x$ ,  $R(x) \subseteq R(w)$ . But  $R(x)$  and  $R(w)$  have the same finite number of members,  $2h+1$ . Thus  $R(x) = R(w)$ . So the midpoint of  $R(x)$  is the midpoint of  $R(w)$ ; that is,  $x = w$ . Hence  $KR(w)$  is true in  $w$  and in no other world.

Now add evidential probabilities to the model as above, with a uniform prior distribution,  $\text{Prob}_{\text{prior}}$ . Since  $KR(w)$  is true in just one of the  $2h+1$  worlds in  $R(w)$ , its evidential probability in  $w$ ,  $\text{Prob}_w(KR(w))$ , is  $1/(2h+1)$ . By increasing the number of positions round the dial while keeping your discriminatory capacities fixed, we can increase  $h$  without limit, and thereby make the evidential probability in  $w$  that one knows  $R(w)$  as small as desired, even though  $R(w)$  is in fact known in  $w$ .

As in section 2, the evidential probability in the model that the proposition is known is not only small, but known to be small, and known to be known to be small, and known to be known to be known to be small, and .... For since  $KR(w)$  is true in only one world, and for any world  $x$   $R(x)$  has  $2h+1$  members,  $\text{Prob}_x(KR(w))$  is always at most  $1/(2h+1)$ . Thus the proposition  $[Pr(KR(w)) \leq 1/(2h+1)]$  is true in every world in the model. Consequently, the proposition  $K^k[Pr(KR(w)) \leq 1/(2h+1)]$  is also true in every world. In other words, one can have any number of iterations of knowledge that the probability of  $R(w)$  is at most  $1/(2h+1)$ .

One cannot avoid these structural results by tightening the conditions for knowledge, short of complete scepticism. For reducing the range of truths known amounts in this setting to increasing the margin for error  $h$ . But, given the symmetry of the situation, the argument holds for any positive margin for error — unless  $h$  is made so large that  $R(w)$  is the whole of  $W$ , which is in effect to say that one learns nothing by looking at the dial.

Even denying the equation of evidence with knowledge would make very little difference to the argument. It would presumably involve postulating one margin for error  $h$  for knowledge and a distinct margin for error  $h^*$  for evidence: the worlds compatible with the total content of one's evidence in  $w$  would be those within a distance  $h^*$  of  $w$ ;  $h^*$  is nonzero too, for more than one position is compatible with one's evidence. That would not affect the argument that  $KR(w)$  is true in no world except  $w$ . Hence the probability on one's evidence in  $w$  of  $KR(w)$  would be  $1/(2h^*+1)$ . By increasing the number of positions for the dial, one can make  $h^*$  arbitrarily high, and therefore the probability on one's evidence in  $w$  that one knows  $R(w)$  arbitrarily low.

We can even generalize the argument from knowledge to rational belief (or justified belief), while keeping an independent standard for evidence (as in the previous paragraph). Unlike knowledge, rational belief is supposed not to require truth. Standardly, epistemic logic for rational belief differs from epistemic logic for knowledge just in replacing the principle that what is known is true by the principle that what it is rational to believe is consistent. For an operator for rational belief (rather than knowledge), a world  $x$  is accessible from a world  $w$  if and only if whatever it is rational for one to believe (rather than whatever one knows) at  $w$  is true at  $x$ . Technically, the

constraint that the accessibility relation be reflexive is relaxed to the constraint that it be *serial*, in the sense that no world has it to no world. The effect is that the T schema  $KA \rightarrow A$  is weakened to the D schema  $JA \rightarrow \neg J\neg A$  (writing  $J$  for ‘it is rational for one to believe that’ or ‘one is justified in believing that’). Of course, dropping the constraint that accessibility be reflexive does not imply adding the constraint that it be non-reflexive. In the present case, since we may build into the example that it is quite clear to one from general background information that one is not suffering from any illusion or systematic distortion of perception; one’s only problem is the limit on one’s powers of perceptual discrimination. Thus, as before, for any worlds  $w, x, y$  and  $z$  in the model, if  $w$  is at least as close to  $x$  as  $y$  is to  $z$  (in number of steps around the circumference) then  $x$  is accessible from  $w$  if  $z$  is accessible from  $y$ . As before, that implies that for some natural number  $h^{**}$  (perhaps distinct from  $h$  and  $h^*$ ), constant across the model, one world is accessible from another if and only they are at most  $h^{**}$  steps apart. In particular, every world is accessible from itself, not by force of any general constraint about rational belief, but simply as a feature of this specific epistemic situation. Rational belief sometimes behaves like knowledge. Thus the structure is just as in the previous paragraph, with the upshot that it can be rational for one to believe a proposition even though it is almost certain on one’s evidence that it is not rational for one believe that proposition.

If the condition for rational belief were relaxed from ‘in all accessible worlds’ to ‘in most accessible worlds’, the inference from  $JA \ \& \ JB$  to  $J(A \ \& \ B)$  would fail. For ‘in most accessible worlds’ will be equivalent to ‘in at least  $k$  accessible worlds’ for some given natural number  $k$  greater than  $h^{**}$  and less than  $2h^{**} + 1$ . Let  $w$  be a world, and let  $A$  be a subset of the set of worlds accessible from  $w$  with exactly  $k$  members, including



the two endpoints  $x$  and  $y$  of that set. Thus  $JA$  is true at  $w$ . We may assume that the total number of worlds is well over  $4h^{**}$ , since otherwise one's eyesight in the model is so bad that when the hand is pointing at twelve o'clock, it is not even rational for one to believe that it is pointing between eleven o'clock and one o'clock. Then from any world other than  $w$ ,  $x$  and  $y$  are not both accessible, so not all members of  $A$  are accessible, so fewer than  $k$  members of  $A$  are accessible, so it is false that most members of  $A$  are accessible. Thus  $w$  is the *only* world at which  $JA$  is true, the only world at which it is rational for one to believe  $A$ . From here the argument proceeds as before.

The foregoing results should still hold on reasonable variations in the prior probability distribution  $R(w)$  that make it slightly non-uniform, for  $KR(w)$  will still be true only in  $w$ , and so its probability (the probability that one is in  $w$ ) will still be low in  $w$  and *a fortiori* everywhere else too. Similarly, making the allowable space of positions for the hand continuous rather than discrete should not make much difference. One would also expect the target phenomenon often to arise in comparable ways when the epistemic accessibility relation  $R$  takes different forms, for example by being linear or multi-dimensional. Nor do psychologically more realistic descriptions of knowledge seem to raise the probability on one's evidence that one knows the strongest relevant proposition one can know, when one does in fact know it. Thus the target epistemic phenomenon seems robust.

Reflection suggests a generalization of the example. One key structural feature of the model is this:

(\*) For all worlds  $w, x$ :  $R(x) \subseteq R(w)$  only if  $x = w$ .

That is, shifting from one world to another (as from  $w$  to  $x$ ) always opens up new epistemic possibilities as well as perhaps closing down old ones. Some worlds are close enough to  $x$  to be epistemically possible in  $x$  but not close enough to  $w$  to be epistemically possible in  $w$ . This is a plausible feature of real-life examples of inexact knowledge. As we move through logical space, our epistemic horizon moves with us. In fact, this more limited feature may do in place of (\*):

(\*\*) Some world  $w$  is such that for all worlds  $x$ :  $R(x) \subseteq R(w)$  only if  $x = w$ .

For, given (\*\*),  $w$  is still the only world in which  $R(w)$  is known, so the evidential probability of  $KR(w)$  will tend to be small in  $w$ , given the subject's imperfect powers of discrimination, even though  $KR(w)$  is always true in  $w$ .

Note that if at least one verifying world  $w$  for (\*\*) has  $R$  to a world other than itself then  $R$  is non-transitive. For suppose that  $Rwx$ ,  $x \neq w$  and  $R(x) \subseteq R(w)$  only if  $x = w$ . Then for some world  $y$ ,  $Rxy$  but not  $Rwy$ , so transitivity fails.

The existence of natural structural generalizations such as (\*) and (\*\*) provides some further confirmation of the robustness of the phenomenon of knowing that is highly improbable on the subject's own evidence.

4. One restrictive feature of the model in section 3 is that the width of the margin for error required for knowledge is in effect treated as beyond doubt, since it is built into the structure of the model. More specifically, since the model has only one world in which

the clock hand has a given position, worlds can differ over what positions are epistemically possible for the hand only by differing over which position it in fact has. Yet it is plausible that there is inexactness in our knowledge of the width of the margin for error in addition to the inexactness in our knowledge of the position of the hand. If so, then in more realistic models the worlds epistemically possible in a given world  $w$  will include some in which the margin for error differs slightly from that in  $w$ , while the position of the hand is the same. In particular, in  $w$  a world  $x$  is epistemically possible in which the margin for error is slightly less than in  $w$ . In such cases we may have  $R(x) \subseteq R(w)$  even though  $x \neq w$ . Pictorially: a sphere may contain a sphere of slightly smaller radius whose centre is a slight distance from the centre of the first sphere. Then whatever is known in  $w$  is also known in  $x$ . In such cases, (\*) and even (\*\*) may fail.<sup>2</sup>

To construct models with a variable margin for error is not hard. But doing so without making *ad hoc* choices is harder. In effect, one must specify higher-order margins for error distinct from the first-order margins for error. There is no obvious non-arbitrary way of determining the relation between the widths of the margins at different orders. By contrast with the simpler case in section 3, it is not clear which models one should be considering.

Nevertheless, in a setting with variable margins for error, one can still give an informal argument for a conclusion similar to that already reached in the case of constant margins. Let  $H(w)$  be the strongest proposition known in  $w$  about the position of the hand (or whatever other non-epistemic fact is relevant). Thus  $H(w)$  may be true at worlds other than  $w$ ; its truth-value remains constant across worlds where the position of the hand is the same, even if the epistemic facts differ. Let  $h$  be the first-order margin for error (the

one relevant to knowledge of the position of the hand) in  $w$ . Thus  $H(w)$  is true in exactly those worlds where the distance of the hand position from that in  $w$  is at most  $h$ . Let  $ME_{<w}$  be true in just those worlds in which the first-order margin of error is less than  $h$ ,  $ME_{>w}$  be true in just those worlds in which the first-order margin for error is greater than  $h$ , and  $ME_{=w}$  be true in just those worlds in which the first-order margin for error is equal to  $h$ . These three possibilities are mutually exclusive and jointly exhaustive. Therefore, by definition of conditional probability:

$$(1) \quad \text{Prob}_w(KH(w)) = \text{Prob}_w(KH(w) \mid ME_{<w}) \cdot \text{Prob}_w(ME_{<w}) + \\ \text{Prob}_w(KH(w) \mid ME_{=w}) \cdot \text{Prob}_w(ME_{=w}) + \\ \text{Prob}_w(KH(w) \mid ME_{>w}) \cdot \text{Prob}_w(ME_{>w})$$

In any world  $x$  in  $ME_{>w}$  some world is epistemically possible in which  $H(w)$  is false, because the first-order margin for error in  $x$  is some  $k > h$ , and a sphere of radius  $k$  cannot be contained in a sphere of radius  $h$ . Thus  $ME_{>w}$  is incompatible with  $KH(w)$ , so  $\text{Prob}_w(KH(w) \mid ME_{>w}) = 0$ . Consequently, (1) simplifies to:

$$(2) \quad \text{Prob}_w(KH(w)) = \text{Prob}_w(KH(w) \mid ME_{<w}) \cdot \text{Prob}_w(ME_{<w}) + \\ \text{Prob}_w(KH(w) \mid ME_{=w}) \cdot \text{Prob}_w(ME_{=w})$$

Since  $\text{Prob}_w(KH(w) \mid ME_{<w}) \leq 1$ , (2) yields:

$$(3) \quad \text{Prob}_w(KH(w)) \leq \text{Prob}_w(ME_{<w}) + \text{Prob}_w(KH(w) \mid ME_{=w}) \cdot \text{Prob}_w(ME_{=w})$$

For simplicity, we may reasonably assume that  $\text{Prob}_w(ME_{<w}) = \text{Prob}_w(ME_{>w})$ , that is, that the first-order margin for error is equally likely to be less or greater than its actual value.

Since  $\text{Prob}_w(ME_{<w}) + \text{Prob}_w(ME_{=w}) + \text{Prob}_w(ME_{>w}) = 1$ ,  $\text{Prob}_w(ME_{<w}) = (1 - \text{Prob}_w(ME_{=w}))/2$ . Therefore, by (3):

$$(4) \quad \text{Prob}_w(KH(w)) \leq (1 - \text{Prob}_w(ME_{=w}))/2 + \text{Prob}_w(KH(w) \mid ME_{=w}) \cdot \text{Prob}_w(ME_{=w})$$

But  $\text{Prob}_w(KH(w) \mid ME_{=w})$  is in effect the probability of  $KH(w)$  in the case considered previously of a constant margin for error ( $h$ ). From that case we have at the very least that  $\text{Prob}_w(KH(w) \mid ME_{=w}) < 1/2$ . Consequently, by (4):

$$(5) \quad \text{Prob}_w(KH(w)) \leq (1 - \text{Prob}_w(ME_{=w}))/2 + \text{Prob}_w(ME_{=w})/2 = 1/2$$

In other words, although you in fact know  $H(w)$  in  $w$ , it is no more probable than not on your evidence in  $w$  that you know  $H(w)$ .

Even if we slightly relax the simplifying assumption that  $\text{Prob}_w(ME_{<w}) = \text{Prob}_w(ME_{>w})$ , the probability on the evidence in  $w$  that  $H(w)$  is known will not rise significantly above evens. Indeed, the probability may well be close to zero. For if the width of the first-order margin for error varies only slightly (as a proportion of  $h$ ) over the worlds epistemically possible in  $w$ , then  $\text{Prob}_w(KH(w) \mid ME_{<w})$  will be close to  $\text{Prob}_w(KH(w) \mid ME_{=w})$ . Therefore, by (2),  $\text{Prob}_w(KH(w))$  will be at most only slightly greater than  $\text{Prob}_w(KH(w) \mid ME_{=w}) \cdot \text{Prob}_w(ME_{<w}) + \text{Prob}_w(KH(w) \mid ME_{=w}) \cdot \text{Prob}_w(ME_{=w}) =$

$\text{Prob}_w(KH(w) \mid ME_{=w}) \cdot \text{Prob}_w(ME_{\leq w})$ . But, as noted above,  $\text{Prob}_w(KH(w) \mid ME_{=w})$  is in effect the probability of  $KH(w)$  in the case already considered of a constant margin for error. That probability goes to zero as the number of hand positions increases. Hence  $\text{Prob}_w(KH(w))$  may well be close to 0 even when the width of the margin for error varies. But even without that stronger conclusion, the result of the informal argument is enough for present purposes. Uncertainty about the width of the margin for error does not undermine the possibility of knowing something without its being probable on one's evidence that one knows it.

5. Examples of the kind considered in previous sections make trouble for accounts of propositional justification on which a sufficient condition for having such justification to believe  $p$  is that the probability of  $p$  on one's evidence exceeds a threshold less than 1. For in such cases one would then have propositional justification to believe the Moore-paradoxical conjunction  $p \ \& \ \neg Kp$ . Consider a world  $w$  in which one knows  $p$  but  $\text{Prob}_w(Kp)$ , the probability on one's evidence in  $w$  that one knows  $p$ , is less than  $1/n$ , so  $\text{Prob}_w(\neg Kp) > (n-1)/n$ . Since one knows  $p$ ,  $\text{Prob}_w(p) = 1$ . By elementary probability theory, it follows that  $\text{Prob}_w(p \ \& \ \neg Kp) > (n-1)/n$ . Thus by letting  $n$  go high enough, we can find a case in which the probability on one's evidence of the Moorean conjunction exceeds the given threshold and the supposedly sufficient condition for justification is met. But it is not plausible that one can have justification to believe a conjunction of the form of "It's raining and I don't know that it's raining". In effect, this point is a sort of abstract generalization of the objection to probabilistic acceptance rules from Moore paradoxes in lottery cases ("My ticket won't win and I don't know that it won't win"),

but it avoids their reliance on the specific assumption that one is not in a position to know that one's ticket won't win.

However, the phenomenon at issue does not merely raise problems for particular philosophical views. More importantly, it raises problems for knowing and acting subjects as such, not just for those knowing and acting subjects who choose to theorize about knowing and acting subjects. Suppose, for instance, that it is rational for Hamlet to perform some action  $A$  if he knows  $p$  and not rational for him to perform  $A$  otherwise (given the relevant background circumstances, such as his other attitudes), and that those facts about rationality are clear to him.<sup>3</sup> Suppose also that although Hamlet knows  $p$ , it is almost certain on his evidence that he does not know  $p$ . Should Hamlet do  $A$ ? Since he knows  $p$ , it is rational for him to do  $A$ . However, since it is almost certain on his evidence that he does not know  $p$ , it is almost certain on his evidence that it is not rational for him to do  $A$ . It is therefore very tempting to say that after all it is not rational for him to do  $A$ . But to say that would be to contradict the conditions of the example. Thus there is pressure to say instead that the example cannot really arise, that if it is almost certain on one's evidence that one does not know  $p$  then one does not really know  $p$ . We have already seen that such pressure should be resisted, since it will not stop short of scepticism. Nevertheless, its effect when we consider particular cases may be to make us withdraw true knowledge ascriptions, under the false impression that they have been defeated by the negative evidence.

Of course, we have no general difficulty with the idea that a claim may be true even though it is almost certain on someone's evidence that it is false — for example, when the claim states the actual outcome of a long sequence of unobserved coin tosses.

What we find harder to accept is the possibility of the same combination when the claim ascribes knowledge to the very person whose evidence is in question at that very time. Their strong evidence that they do not know  $p$  seems incompatible with whatever sort of reliance on  $p$  is mandated by their knowing  $p$ . But the apparent incompatibility is an illusion, which can seriously distort our assessment of particular ascriptions of knowledge and with it our epistemological theorizing.

The phenomenon that we find hard to accept is an extreme case of anti-luminosity, the failure of non-trivial states to satisfy the constraint that whenever one is in them one is in a position to know that one is in them (Williamson 2000: 93-113). In evidential terms, when luminosity fails one is in a state  $S$  even though it is not certain on one's evidence that one is in  $S$ . In the present case,  $S$  is the state of knowing  $p$ . But the phenomenon at issue is stronger than a mere failure of luminosity, since it involves being in  $S$  even though it is almost certain on one's evidence that one is *not* in  $S$ .

If we do resist the sceptical pressure, and acknowledge the possibility of the phenomenon, then we must regard Hamlet's problem above as a genuine problem for him, an instance of a more general practical problem for agents acting on inexact knowledge, not as something to be dissolved by epistemological redescription. The solution is clear from the conditions of the example: it is rational for Hamlet to do  $A$  if and only if he knows  $p$ , and he does know  $p$ , so it *is* rational for him to do  $A$ . The trouble is that he has strongly misleading evidence about those facts. But that does not mean that they are not facts after all; it just means that he has a good excuse — not justification — for not doing  $A$ .<sup>4</sup>



By a parallel argument, just as it can be rational for one to do  $A$  even though it is extremely improbable on one's evidence that it is rational for one to do  $A$ , so it can happen that one *should* do  $A$ , even though it is extremely improbable on one's evidence that one should do  $A$ .

We must check these general comments by examining a specific case in more detail. That is the business of the following final section.

6. Competent deduction is a way of extending our knowledge, as in mathematical reasoning. A natural formulation of the underlying principle is this form of *multi-premise closure*:

MPC If one believes a conclusion on the basis of competent deduction from premises each of which one knows, one knows the conclusion.

For present purposes we can leave the phrase 'on the basis of competent deduction' vague, and assume that something in the spirit of MPC is correct.<sup>5</sup>

MPC faces a strong challenge from cases of reasoning from many premises, each with a small independent risk of error, where those risks combine into a large risk of error for the conclusion. Such situations arise most simply when the conclusion is the conjunction of the premises, since then the conclusion entails the premises and any error in a premise involves an error in the conclusion; the deduction of the conclusion from the premises also takes a particularly elementary form. Suppose that each premise is true, so the conclusion is true too. We may further suppose that, according to ordinary standards,

one knows each premise. For example, each premise may state a separate matter of particular historical fact that one has carefully checked. To deny that one knows a given premise would look like extreme scepticism. From those propositions one has competently deduced their conjunction and believes it on that basis. By MPC, one thereby knows the conclusion. Nevertheless, much past experience may show that there is a small non-zero rate of error for such carefully checked historical claims, on the basis of which it is almost inevitable that the long conjunction contains several false conjuncts. Thus one's belief in the conjunction may seem to have too strong a risk of error to constitute knowledge. This is a variant of the Preface Paradox, adapted to knowledge.

On a fallibilist conception of knowledge, it seems, the risk of error for any one premise may be within the threshold for knowledge, while the risk of error for the conclusion is outside the threshold; thus MPC fails. By contrast, on an infallibilist conception, knowledge requires that there be no risk of error, rather than at most a small risk, so MPC may hold; but such infallibilism seems to lead to scepticism. Alternatively, one might try to preserve MPC by postulating some sort of variation across contexts of utterance in the reference of 'know', or some loss of knowledge in the premises through the very act of deducing the conclusion.<sup>6</sup>

Once we recognize the phenomenon of knowing when it is almost certain on one's evidence that one fails to know, we can see a possible diagnosis of the problem cases for MPC which allows us to keep MPC while neither falling into scepticism nor postulating truth-conditions for 'knowledge'-ascriptions of any non-standard kind. One does indeed know each premise, without knowing that one knows it. Since one believes the conclusion on the basis of competent deduction from the premises, by MPC one also

knows the conclusion, although without knowing that one knows it. For each premise, it is very probable on one's evidence that one knows it. However, it is very improbable on one's evidence that one knows every premise. Given that one knows the conclusion (the conjunction) only if one knows every premise, it is very improbable on one's evidence that one knows the conclusion. Since we are tempted to conceive knowing as at least somewhat luminous, we are tempted to deny that one knows the conclusion.<sup>7</sup>

We can give a formal epistemic model of that description. As noted in section 1, such a model automatically validates not just MPC but logical omniscience, an unrealistically strong version of multi-premise closure that does not even require the agent to have carried out the relevant deductions, competently or otherwise. As a positive argument for MPC, it might well be accused of both begging the question and proving too much. However, its role here is different. It is being used to defuse an objection to MPC, by showing that even on the assumption of the strongest possible version of multi-premise closure, one would predict the occurrence of epistemic phenomena that it is very tempting to misinterpret along the lines of the objection as counter-examples to multi-premise closure. In particular, if we treat ascriptions of knowledge as defeated by a low probability of knowing, but not by a high probability short of 1 of knowing, on the subject's evidence, then we shall tend to judge that the subject knows each conjunct without knowing the conjunction, even though the conditions for MPC are satisfied; we are deceived by the false appearance of a counterexample to MPC.<sup>8</sup>

Here are the details of such a model.<sup>9</sup> For worlds we use  $n$ -tuples of numbers drawn from the set  $\{0, 1, \dots, 2k\}$ , where  $n$  is the number of premises (conjuncts) and  $k$  is a large natural number. Thus there are  $(2k+1)^n$  worlds. The  $n$  components of a world

represent its locations on  $n$  independent dimensions of a state space. The  $i$ th dimension is the one relevant to the  $i$ th premise  $p_i$ . Let the  $i$ th component of the  $n$ -tuple  $w$  be  $w_i$ , and the world just like  $w$  except that its  $i$ th component is  $m$  be  $w[i|m]$ , so  $w[i|m]_i = m$  and  $w[i|m]_j = w_j$  if  $i \neq j$ . Let  $p_i$  be true in  $w$  if and only if  $w_i > 0$ . Let  $x$  be epistemically possible in  $w$  ( $Rwx$ ) if and only if for all  $i$ ,  $|w_i - x_i| \leq k$ , that is,  $w$  and  $x$  do not differ by ‘too much’ in any of their respective components. In effect, a margin for error is applied to each of the  $n$  dimensions separately. The relation  $R$  is obviously reflexive and symmetric. We can easily check that for any world  $w$ ,  $p_i$  is known ( $Kp_i$  is true) in  $w$  if and only if  $w_i > k$ .<sup>10</sup> Similarly, we can check that  $p_i$  is known to be known ( $KKp_i$  is true) in  $w$  if and only if  $w_i > 2k$ ; hence  $p_i$  is not known to be known ( $KKp_i$  is not true) in any world in this model. In particular, in the world  $\langle 2k, \dots, 2k \rangle$ , each premise  $p_i$  is known and none is known to be known.

As usual, the prior probability distribution is uniform: for any world  $w$ ,  $\text{Prob}_{\text{prior}}(\{w\}) = 1/(2k+1)^n$ . We must check that the model makes the  $n$  dimensions probabilistically independent. For any given  $i$ , a proposition  $q$  is  $i$ -based if and only if for all worlds  $x$  and  $y$ , if  $x_i = y_i$  then  $q$  is true in  $x$  if and only if  $q$  is true in  $y$  ( $1 \leq i \leq n$ ). That is, whether an  $i$ -based proposition is true in a world depends only on the  $i$ th component of that world. In particular,  $p_i$  is an  $i$ -based proposition. Obviously, the negation of any  $i$ -based proposition is also  $i$ -based, as is any conjunction of  $i$ -based propositions. We can also prove that whenever  $q$  is an  $i$ -based proposition, so is  $Kq$ .<sup>11</sup> Thus  $Kp_i$  and  $KKp_i$  are also  $i$ -based propositions. Then we can prove that whenever for each  $i$   $q_i$  is an  $i$ -based proposition,  $q_1, \dots, q_n$  are mutually probabilistically independent on the evidence in any world  $w$ , in the usual sense that the probability (on the evidence in  $w$ ) of their conjunction

is the product of the probabilities (on the evidence in  $w$ ) of the conjuncts.<sup>12</sup> Although a model could have been constructed in which the evidence at some worlds establishes epistemic interdependences between the different dimensions, for present purposes we can do without such complications. In particular,  $p_1, \dots, p_n$  are mutually probabilistically independent on the evidence in any world, as are  $Kp_1, \dots, Kp_n$ . But, on the evidence in the world  $\langle 2k, \dots, 2k \rangle$ , for any given  $i$ , the probability that  $p_i$  is known is  $k/(k+1)$ .<sup>13</sup> By probabilistic independence, the probability of the conjunction  $Kp_1 \cap \dots \cap Kp_n$  is  $(k/(k+1))^n$ . That is the probability that each conjunct is known. But, by the logical omniscience built into the model, knowing a conjunction ( $K(p_1 \cap \dots \cap p_n)$ ) is equivalent to knowing each conjunct. Thus the probability on the evidence in  $\langle 2k, \dots, 2k \rangle$  that the conjunction  $p_1 \cap \dots \cap p_n$  is known is also  $(k/(k+1))^n$ . For fixed  $k$ , this probability becomes arbitrarily close to 0 as  $n$  becomes arbitrarily large. Thus, for suitable  $k$  and  $n$ , the world  $\langle 2k, \dots, 2k \rangle$  exemplifies just the situation informally sketched: for each conjunct one knows it without knowing that one knows it, and it is almost but not quite certain on one's evidence that one knows the conjunct; one also knows the conjunction without knowing that one knows it, and it is almost but not quite certain on one's evidence that one does *not* know the conjunction.

Of course, in some examples one's epistemic position with respect to each conjunct is better: one not only knows it but knows that one knows it. If one also knows the relevant closure principle, and knows that one satisfies the conditions for its application, one may even know that one knows the conjunction. Consequently, the probability on one's evidence that one knows the conjunction is 1. However, the previous pattern may still be repeated at a higher level of iterations of knowledge. For example, for

each conjunct one knows that one knows it without knowing that one knows that one knows it, and it is almost but not quite certain on one's evidence that one knows that one knows the conjunct; one also knows that one knows the conjunction without knowing that one knows that one knows it, and it is almost but not quite certain on one's evidence that one does *not* know that one knows the conjunction. To adapt the previous model to this case, we can simply expand the set of worlds by using  $n$ -tuples of numbers from the set  $\{0, 1, \dots, 3k\}$  rather than  $\{0, 1, \dots, 2k\}$ , leaving the definitions of the epistemic possibility relation  $R$  and the truth-conditions of the  $p_i$  unchanged (so  $p_i$  is true in  $w$  if and only if  $w_i > 0$ ); then  $\langle 3k, \dots, 3k \rangle$  is a world of the required type. More generally, if one uses as worlds  $n$ -tuples of numbers from the set  $\{0, 1, \dots, hk\}$ , leaving the other features of the model unchanged, then  $\langle hk, \dots, hk \rangle$  will be a world at which one has  $h-1$  but not  $h$  iterations of knowledge of each conjunct, and it is almost but not quite certain on one's evidence that one has  $h-1$  iterations of knowledge of the conjunct; one also has  $h-1$  but not  $h$  iterations of knowledge of the conjunction, and it is almost but not quite certain on one's evidence that one does *not* have  $h-1$  iterations of knowledge of the conjunction.<sup>14,15</sup>

Many other variations can be played on the same theme. The general idea is this. Suppose that the epistemic status  $E$  satisfies an appropriate principle of multi-premise closure. In some situations, one attains  $E$  with respect to each conjunct, without knowing that one does so (this is possible by the anti-luminosity argument). By multi-premise closure, one also attains status  $E$  with respect to the conjunction, without knowing that one does so. Then for each conjunct it may be almost certain on one's evidence that one attains  $E$  with respect to it, even though it is almost certain on one's evidence that one does *not* attain  $E$  with respect to the conjunction. If we treat ascriptions of  $E$  as defeated

by a low probability of E, but not by a high probability short of 1 of E, on the subject's evidence, then we shall tend to judge that the subject attains E with respect to each conjunct but not with respect to the conjunction, even though the conditions for multiple-premise closure principle are satisfied; we are deceived by the false appearance of a counterexample to the multi-premise closure principle.

7       The considerations of this paper raise a more general question. Knowledge claims are often thought to be defeated by various sorts of misleading evidence. In how many cases is the correct account that the subject knows, even though it is almost certain on the subject's evidence at the time that they do not know? That is left as an open question.

## Notes

\* Despite its inclusion in a volume mainly devoted to Earl Conee and Richard Feldman's influential work on evidentialism, this paper does not react to their ideas. Rather, it illustrates how the systematic investigation of evidence can play a central role in a quite different approach to epistemology. The main idea for the paper arose in response to discussion with John Hawthorne. Earlier versions of the material were presented in classes and a discussion groups at Oxford, at an Arché workshop on Basic Knowledge at St Andrews, where Philip Ebert was the commentator, at a workshop on epistemology at the University of Wisconsin at Madison, where Jim Pryor was the commentator, at graduate conferences at University College London and the University of Miami, at a workshop on mathematical methods in philosophy in Banff (Alberta), at a symposium for Jaakko Hintikka in Copenhagen, at the Moral Sciences Club in Cambridge, the Argentine Society for Analytic Philosophy in Buenos Aires and the Serbian Philosophical Society in Belgrade, and at colloquia in the philosophy departments of Lund University, Nottingham University, the University of Texas at Austin and Brandeis University. I thank all those who have contributed to the paper with their comments and questions.

1 For an argument that epistemically possible propositions can have probability 0, even when infinitesimal probabilities are allowed, see Williamson 2007b.



2 On an epistemic account of vagueness, such variable margins for error yield distinctive forms of higher-order vagueness. Williamson (1999: 136-8) argues that if the ‘clearly’ operator for vagueness obeys the analogue of the B (for ‘Brouwersche’) axiom  $p \rightarrow K\neg K\neg p$  (which corresponds to the condition of symmetry on  $R$ ) then any formula with second-order vagueness has  $n$ th-order vagueness for every  $n > 2$ , but does not endorse the B axiom for ‘clearly’. In response, Mahtani (2008) uses variable margins for error to argue against the B axiom for ‘clearly’ and suggests that they allow vagueness to cut out at any order. Dorr (2008) provides a formal model in which he proves Mahtani’s suggestion to hold. These arguments all have analogues for the overtly epistemic case.

3 For a recent discussion of the relation between knowledge and reasons for action see Hawthorne and Stanley 2008. By what was argued in section 3, we can make a corresponding argument for rational belief in place of knowledge.

4 For a critique of internalism about justification see Williamson 2007a.

5 Such a principle is called ‘intuitive closure’ at Williamson 2000: 117-18.

6 For discussion of MPC in relation to the Preface Paradox see Hawthorne 2004: 46-50, 154, 182-3. Similar cases involving future chances are used in Hawthorne and Lasonen-Aarnio 2009 against the safety conception of knowledge in Williamson 2000; for a reply see Williamson 2009.

7 A similar problem arises for single-premise closure principles when one competently carries out a long chain of single-premise deductive steps, each with a small epistemic probability of inferential error (in the multi-premise case, for simplicity, one's deductive competence is treated as beyond doubt); see Lasonen-Aarnio 2008 for discussion. Here is a parallel account of that case. One knows the premise without knowing that one knows it. For each deductive step, one carries it out competently without knowing that one does so. By single-premise closure, one knows the conclusion, without knowing that one knows it. For each deductive step, it is very probable on one's evidence that one carries it out competently. However, it is very improbable on one's evidence that one carries out every deductive step competently. Since it is granted that one knows the conclusion only if one carries out every deductive step competently, it is very improbable on one's evidence that one knows the conclusion.

8 The usual form of epistemic modelling is not appropriate for treating possible errors in deductive reasoning, since logical omniscience suppresses the dependence of inferential knowledge on correct inferential processes.

9 The technical details are taken from Williamson 2009.

10 Proof: Suppose that  $w_i > k$ . If  $Rwx$  then  $|w_i - x_i| \leq k$ , so  $x_i > 0$ , so  $x \in p_i$ . Thus  $w \in Kp_i$ . Conversely, suppose that  $w_i \leq k$ . Then  $Rww[i|0]$ , for  $|w_i - w[i|0]_i| = |w_i - 0| = w_i \leq k$  and if  $i \neq j$  then  $|w_j - w[i|0]_j| = 0$ ; but  $w[i|0] \notin p_i$  because  $w[i|0]_i = 0$ , so  $w \notin Kp_i$ .

11 Proof: Suppose that  $q$  is  $i$ -based and  $x_i = y_i$ . Suppose also that  $x \notin Kq$ . Then for some  $z$ ,  $Rxz$  and  $z \notin q$ . But then  $Ryy[i|z_i]$ , for  $|y_i - y[i|z_i]| = |y_i - z_i| = |x_i - z_i|$  (because  $x_i = y_i$ )  $\leq k$  (because  $Rxz$ ), and if  $i \neq j$  then  $|y_j - y[i|z_i]| = 0$ . Moreover,  $y[i|z_i] \notin q$  because  $z \notin q$ ,  $q$  is  $i$ -based and  $y[i|z_i]_i = z_i$ . Hence  $y \notin Kq$ . Thus if  $y \in Kq$  then  $x \in Kq$ . By parity of reasoning the converse holds too.

12 Proof: Set  $\#(i, q, w) = \{j: 0 \leq j \leq 2k, w[i|j] \in q \text{ and } |w_i - j| \leq k\}$  for any  $w \in W$ ,  $q \subseteq W$ ,  $1 \leq i \leq n$ . For each  $i$ , let  $q_i$  be  $i$ -based. Let  $\cap q_i = q_1 \cap \dots \cap q_n$ . For  $w \in W$ ,  $R(w) \cap \cap q_i = \{x: \forall i, x_i \in \#(i, q_i, w)\}$ , since for each  $i$  and  $x \in W$ ,  $x \in q_i$  iff  $w[i|x_i] \in q_i$  since  $q_i$  is  $i$ -based. Since  $\text{Prob}_{\text{prior}}$  is uniform,  $\text{Prob}_w(\cap q_i) = |\{R(w) \cap \cap q_i\}|/|R(w)|$  for  $w \in W$ . But  $|R(w) \cap \cap q_i| = |\{x: \forall i, x_i \in \#(i, q_i, w)\}| = |\#(1, q_1, w)| \dots |\#(n, q_n, w)|$ . By the special case of this equation in which each  $q_i$  is replaced by  $W$  (which is trivially  $i$ -based for any  $i$ ),  $|R(w)| = |\#(1, W, w)| \dots |\#(n, W, w)|$ . Consequently:

$$\text{Prob}_w(\cap q_i) = (|\#(1, q_1, w)| \dots |\#(n, q_n, w)|) / (|\#(1, W, w)| \dots |\#(n, W, w)|).$$

For any given  $i$ , consider another special case in which  $q_j$  is replaced by  $W$  whenever  $i \neq j$ . Since  $n-1$  of the ratios cancel out,  $\text{Pr}_w(q_i) = |\#(i, q_i, w)|/|\#(i, W, w)|$ . Therefore  $\text{Prob}_w(\cap q_i) = \text{Prob}_w(q_1) \dots \text{Prob}_w(q_n)$ , as required.

13 Proof: We have already established that  $x \in Kp_i$  iff  $x_i > k$ . Thus, in the notation of the previous footnote,  $\#(i, Kp_i, \langle 2k, \dots, 2k \rangle) = \{j: k < j \leq 2k\}$ , so  $|\#(i, Kp_i, \langle 2k, \dots, 2k \rangle)| = k$ , while  $\#(i, W, \langle 2k, \dots, 2k \rangle) = \{j: k \leq j \leq 2k\}$ , so  $|\#(i, W, \langle 2k, \dots, 2k \rangle)| = k+1$ . By the formula for  $\text{Prob}_w(q_i)$  in the previous footnote (with  $Kp_i$  in place of  $q_i$ ),  $\text{Prob}_{\langle 2k, \dots, 2k \rangle}(Kp_i) = k/(k+1)$ .

14 A similar generalization to higher iterations of knowledge is possible for the case of multiple risks of inferential error in a single-premise deduction. One has at least  $n$  iterations of knowledge of the premise. For each deductive step, one has  $n-1$  but not  $n$  iterations of knowledge that one carried it out competently. By single-premise closure and plausible background assumptions, one has  $n$  but not  $n+1$  iterations of knowledge of the conclusion. For each deductive step, it is very probable on one's evidence that one has at least  $n-1$  iterations of knowledge that one carried it out competently. However, it is very improbable on one's evidence that one has at least  $n-1$  iterations of knowledge that one carried out every deductive step competently. Since it is granted that one has at least  $n$  iterations of knowledge of the conclusion only if one has at least  $n-1$  iterations of knowledge that one carried out every deductive step competently, it is very improbable on one's evidence that one has at least  $n$  iterations of knowledge of the conclusion.

15 See Williamson 2008 for more discussion of the structure and semantics of higher-order evidential probabilities. The phenomenon discussed in the text involves the apparent loss of only one iteration of knowledge between premises and conclusion. However, the apparent absence of a given number of iterations of knowledge can cause doubts about all lower numbers of iterations, by a domino effect, since lack of knowledge that one has  $n+1$  iterations implies lack of warrant to assert that one has  $n$  iterations (Williamson 2005: 233-4).

## Bibliography

- Dorr, C. 2008. 'How vagueness could cut out at any order'. MS.
- Greenough, P., and Pritchard, D., eds. 2009. *Williamson on Knowledge*, Oxford: Oxford University Press, forthcoming.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Clarendon Press.
- Hawthorne, J., and Lasonen-Aarnio, M. 2009. 'Knowledge and objective chance', in Greenough and Pritchard.
- Hawthorne, J., and Stanley, J. 2008. 'Knowledge and action'. *The Journal of Philosophy* 105: 000-000.
- Hintikka, J. 1962. *Knowledge and Belief*. Ithaca, N.Y.: Cornell University Press.
- Lasonen-Aarnio, M. 2008. 'Single premise deduction and risk'. *Philosophical Studies* 141: 157-73.
- Lemmon, E. J. 1967. 'If I know, do I know that I know?'. In A. Stroll (ed.), *Epistemology*. New York: Harper & Row.
- Mahtani, A. 2008. 'Can vagueness cut out at any order?'. *Australasian Journal of Philosophy* 86: 499-508.
- Radford, C. 1966. 'Knowledge — by examples'. *Analysis* 27: 1-11.
- Williamson, T. 1999. 'On the structure of higher-order vagueness'. *Mind* 108: 127-43.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williamson, T. 2005. 'Contextualism, subject-sensitive invariantism and knowledge of knowledge'. *The Philosophical Quarterly* 55: 213-35.
- Williamson, T. 2007a. 'On being justified in one's head'. In M. Timmons, J. Greco and

- A. Mele (eds.), *Rationality and the Good: Critical Essays on the Ethics and Epistemology of Robert Audi*. Oxford: Oxford University Press.
- Williamson, T. 2007b. 'How probable is an infinite sequence of heads?' *Analysis* 67: 173-80.
- Williamson, T. 2008. 'Why epistemology can't be operationalized.' In Q. Smith (ed.), *Epistemology: New Philosophical Essays*. Oxford: Oxford University Press.
- Williamson, T. 2009. 'Reply to John Hawthorne and Maria Lasonen-Aarnio', in Greenough and Pritchard.