


At the threshold of knowledge

Daniel Rothschild¹ · Levi Spectre² 

© Springer Science+Business Media Dordrecht 2017

Abstract We explore consequences of the view that to know a proposition your rational credence in the proposition must exceed a certain threshold. In other words, to know something you must have evidence that makes rational a high credence in it. We relate such a threshold view to Dorr et al.’s (Philosophical Studies 170(2):277–287, 2014) argument against the principle they call FAIR COINS: “If you know a coin won’t land tails, then you know it won’t be flipped.” They argue for rejecting FAIR COINS because it leads to a pervasive skepticism about knowledge of the future. We argue that the threshold view of evidence and knowledge gives independent grounds to reject FAIR COINS.

Keywords Evidence · Knowledge · Thresholds · Tolerance principle · Probability · Rational-belief · Inference · Epistemological skepticism about the future

To know something requires it be supported by one’s evidence. One natural way of modeling the support an agent’s total evidence gives to a proposition is to use her rational credence, a real number between 0 and 1 that represents the degree to which she should be confident in the proposition’s truth given her evidence.¹ The requirement that knowledge of a proposition be supported by evidence, on this

¹ We need not assume that there is only one degree of confidence that all rational agents must have given the same evidence. But in the present setting, the probabilities will be straightforward. We will therefore refrain from defining a credence function and move freely between probabilities, credences, and degrees of confidence.

✉ Levi Spectre
levi.spectre@gmail.com

¹ University College London, London, United Kingdom

² The Open University of Israel, Ra’anana, Israel

framework, could amount to a requirement that the evidence for a given proposition allow a rational agent to have a credence that is greater than some number k . In non-trivial threshold views k is strictly less than 1.

Threshold views, in this sense, only put *requirements* on knowledge, rather than defining knowledge: passing the threshold is necessary but not sufficient for having knowledge. Threshold views have surprising implications, however. One well-known consequence is the potential failure of the closure of knowledge under conjunction introduction (and other multi-premise logical operations). If there are probabilistically independent propositions p and q each of which has a probability on the evidence just above the threshold k , then the probability of their conjunction will be below the threshold, and thus while one might be in a position to know p and to know q one isn't in a position to know $p \& q$.

A less explored implication of the threshold view—one focus of this paper—is that there can be cases where one knows a disjunction of two propositions p and q where p has a very high probability (but one just below k) and q has a low probability. In this case while one might not be in a position to know the very probable p one can know the disjunction $p \vee q$, whose probability while still less than one is, nonetheless, above the threshold k .

We assess an ingenious argument of Cian Dorr et al. (2014, henceforth, DGH) from the perspective of the threshold view. Their argument is against a principle they call FAIR COINS, which has the following form “If you know that a coin is fair, and for all you know it is going to be flipped, then for all you know it will land tails.” (DGH, p. 277) This is an intuitively plausible principle: it seems obvious that the only way to know that a coin won't land a certain way is to know that it won't be flipped at all. DGH give an intriguing example that aims to show that one must reject FAIR COINS in order to avoid a pervasive skepticism about the future.

We argue here that from the perspective of the threshold view FAIR COINS inevitably fails. This means that there is a reason to reject FAIR COINS independent of an aversion to skepticism. We show also that DGH's argument takes the form of a sorites argument, but one whose inductive premise (closely related to FAIR COINS) is not plausible from the perspective of the threshold view.

Our goal is not to argue for the threshold view of knowledge, but to show that on it one can easily resolve the problems DGH present using it. Other views of the evidential requirements of knowledge might also be used to reject FAIR COINS in a way similar to that which we will argue for here.² The threshold view of the evidential requirements of knowledge, though, is tractable and appealing, and thus provides a good model for seeing what might be wrong with the intuitively appealing FAIR COINS principle.³

² Moreover, though we assume a non-maximal threshold, we are not arguing here that a probability 1 threshold cannot be used to the same end that we use ours here. But see Bacon (2014) for an argument that a probability 1 conception of knowledge leads to trouble in the type of case DGH advance. And see footnote 16 below.

³ Note however, that in Rothschild and Spectre (2016), we give another problematic case based on DGH's basic setup which the threshold view does not help resolve.

Section 1 presents DGH's argument against FAIR COINS. Section 2 shows how the threshold view motivates rejection of FAIR COINS independently of any distaste for skepticism about the future. The section also explains why on the face of things FAIR COINS is so compelling. We conclude in Sect. 3 with a few remarks about knowing disjunctions on the threshold views. We there also draw an analogy with the combination of multi-premise closure and below probability 1 thresholds (as in the Lottery and Preface paradoxes) that we find instructive.

1 DGH

The argument against FAIR COINS has two major steps: First, DGH put forward a scenario that shows the inconsistency of FAIR COINS with knowledge that an extremely unlikely (but higher than probability 0) sequence of coin tosses will not take place. Second, they connect this inconsistency with skepticism about the future: FAIR COINS-like principles—their analogy is a principle they label AUTUMN LEAF (DGH, p. 279)—if valid, will render unknowable, many things about the future that we take agents to regularly know. Their conclusion, then, is that if we want to avoid a commitment to pervasive skepticism about the future—and DGH surely do—we must reject FAIR COINS and its analogs.

Yet because FAIR COINS is so intuitive, and due to the fact that not any knowledge can rely on something simply having a high probability, it seems that DGH have provided the skeptic with an attractive argument: Accept the highly intuitive FAIR COINS and given DGH's step two (a method of devising principles with the same upshot as FAIR COINS), you are a skeptic about the future. Our argument forestalls this kind of argumentative move. Without questioning the two major steps of DGH's argument, we argue that there is solid reason, independent of any (theoretically unmotivated?) distaste for skepticism, to reject FAIR COINS. Knowledge or rational belief thresholds.

That will be the major aim of the next section. We start here with the first stage of DGH's argument:

1000 FAIR COINS are laid out one after another: $C_1, C_2, \dots, C_{1000}$. A coin flipper will flip the coins in sequence until either one lands heads or they have all been flipped. Then he will flip no more. You know that this is the setup, and you know everything you are in a position to know about which coins will be flipped and how they will land. In fact, C_2 will land heads, so almost all of the coins will never be flipped. In this situation it is plausible that, before any of the coins are flipped, you know that C_{1000} will not be flipped—after all, given the setup, C_{1000} will be flipped only in the bizarre event that the previous 999 fair coins all land tails. [...] We can regiment the puzzle as an inconsistent tetrad:

- (1) You know that C_{1000} will not be flipped.
- (2) For each coin C_n : If you know that C_n will not be flipped, then you know that C_{n-1} will not land tails.
- (3) For each coin C_n : If you know that C_n will not land tails, then you know that C_n will not be flipped. [FAIR COINS]

- (4) You don't know that C_1 will not be flipped. The contradiction is obvious: the negation of (4) follows from (1–3) by a long sequence of modus ponens. But (4) is obviously true, since C_2 will be flipped. (2) is hard to deny, given that you know the setup. Assuming the anti-skeptical (1), we are therefore forced to deny (3). So, for some coin, you know that it won't land tails, even though you don't know—and are not in a position to know—that it won't be flipped.

(DGH, pp. 278–279)

There is a temptation here to deny (1) rather than accept that you can know that a coin will not land tails when you don't know whether it will or will not be flipped. And the temptation seems well founded: we often think that no matter how improbable a lottery loss is, you don't know that you've lost until you have at least some non-purely probabilistic access to its result. Simply by believing a highly probable true proposition, an agent will intuitively often fail to know it. So in order to run the argument against their target, DGH need to show that the tempting denial of (1) carries consequences that are worse than the denial of FAIR COINS.

The consequence they draw, is that if FAIR COINS is valid, we must accept a pervasive skepticism about the future. And since lots of knowledge about the future involves at least some low chance condition that must not take place to allow for knowledge—this is roughly the second stage of DGH's argument (pp. 279–280)—the coin toss scenario can be converted into a multitude of cases that regard everyday knowledge about the future, knowledge we intuitively have. Here DGH rely on an analogy between the chancy physical processes that determine future states of affairs and a series of coin tosses. They present a stylized case of an autumn leaf that has a certain probability of falling off a tree in each 1 h period. Their idea is that you know that by the end of winter the leaf will fall off even if in each period its chance of falling is the same as the chance of a (potentially biased) coin landing heads. While this analogy has some force it is not that clear to us how much of our knowledge of the future does rely on processes that share a sufficiently similar underlying probabilistic structure to that of a series of coin tosses.⁴ If we accept that we do not know that a lottery ticket with a low probability will lose, then we must accept that knowledge of high probability future events is sensitive to the structure of the chancy mechanisms underlying them. It is not so clear just how high the cost of denying (1) is.

2 K-THRESHOLD

The inconsistency that DGH derive depends, on the one hand, on their anti-skeptical assumption—that you know C_{1000} won't be flipped. This assumption entails that “there is a smallest number n such that you know that C_n will not be flipped” (DGH,

⁴ Their leaf example is highly unrealistic: the obvious physical parallel is that of radioactive decay. Thus, if we embrace FAIR COINS we might have to accept that we cannot *know* that a set of unstable particles will decay in any finite amount of time.

p. 278). On the other hand, premises (2) and FAIR COINS (3), contradict the assumption. They entail that there is no “smallest number n such that you know that C_n will not be flipped.” This is because together (2) and (3) entail:

- (5) For each coin C_n : If you know that C_n will not be flipped, then you know that C_{n-1} will not be flipped.

By contraposition,

- (6) For each coin C_n : If you don't know that C_{n-1} will not be flipped, you don't know that C_n will not be flipped.

(5) and (6) say that two beliefs relating to neighboring coins, n and $n - 1$, will either both amount to knowledge, or will both fall short of knowledge. This means that a difference in probability between the possibility of a sequence of n tails tosses and $n - 1$ tails tosses isn't enough to make a difference in your knowledge state. Like “the heap” we can start from your assumed knowledge and go back to what is a clear case of ignorance—from a heap to 0 grains—or we can go forward from your ignorance to where clearly you know—from 1 grain to a heap.⁵ But unlike the relation between grains and heaps, in the case of knowledge, these tolerance type principles enjoy little plausibility.⁶ One notable and natural reason to think (5) and (6) are false is:

k-THRESHOLD: One can know the truth of statement A, only if on the evidence A's probability is greater or equal to k (where k is a real number between 0 and 1).

Perhaps there are those who would want to put forward a different knowledge condition that falsifies (5) and (6). Without denying that this can be done, for three reasons *k*-THRESHOLD is best suited for our purposes here. First, it is probably the most prevalent non-tolerance type knowledge condition among theorists. Among them are theorists implicitly accept a version of *k*-THRESHOLD by a combined commitment to the theses that knowledge entails rational (or justified) belief and that the later has a non-maximal sharp threshold (the Lockean thesis). This prevalence is beneficial because it will help us show that at least implicitly FAIR COINS is already widely rejected. Second and relatedly, *k*-THRESHOLD isn't, at least not directly and exclusively, a condition on knowledge that is accepted in order to avoid skepticism. So in putting it forward we are not begging any skeptical question.⁷ But

⁵ Thanks here to Matti Eklund.

⁶ It will come as little surprise that on an understanding of knowledge as an absolute graded notion like “flat”, it will be non-trivial to evade skeptical conclusions. A classical argument to this effect is Unger (1975; pp. 65–68) that almost nothing is flat. Analogies Unger draws with epistemic terms lead him to the conclusion that we have almost no knowledge. Both Lewis (1979) and Dretske (1981) keep the analogies but draw conclusions that avoid skepticism. See Cohen (1987) for a more detailed discussion. The assumptions that need to be denied within DGHs example, as we show in this section, bear more than an artificial similarity with this dialectic. The assumption is that there is some tolerance principle that governs knowledge (and rational belief).

⁷ A skeptic about the future can even accept *k*-THRESHOLD because she accepts it in other knowledge contexts. She could accept, in other words, that had you been able to know that C_n will not be flipped, it

mainly, third, k -THRESHOLD leads to the failure of FAIR COINS and can help locate the reason it is misleadingly compelling. Or so we will now argue.

We start by arguing that k -THRESHOLD⁸ provides direct reason to reject (5) and (6). The basic reason is simple. For any n the probability that C_n will not be flipped is greater than the probability that C_{n-1} will not be flipped. So, if there is a probabilistic threshold on knowledge it makes sense that one can know C_n will not be flipped without knowing that C_{n-1} will not be flipped.

Of course (5) and (6), the inductive principles for DGH's sorites, are not themselves the same as the FAIR COINS assumption, but rather just the inductive premises. It will be useful to show how k -THRESHOLD is actually in tension with FAIR COINS itself. As before, let n be the first coin such that you know

(7) C_n will not be flipped.

We, then, assume that (a) coin n is the first coin that satisfies the k -THRESHOLD, and (b) the setup is known for certain.

Assumption (a) amounts to the claim that while your evidence makes the probability of C_n not being flipped $\geq k$, the probability of C_{n-1} not being flipped is $< k$. Why should this be the case? Your evidence about C_n , is purely statistical: you do not know C_n will not be flipped by perception, testimony or other routes. On the assumption that you *do* know C_n will not be flipped then there is no reason why that statistical evidence does not amount to knowledge in this case.⁹ The only relevant difference between the proposition that C_{n-1} will not be flipped and the proposition

Footnote 7 continued

would not follow that you knew that C_{n-1} will not be flipped (for any n). On this type of skepticism, knowledge about the future could fail (almost) universally because some other necessary condition on knowledge is not satisfied. For instance, we can imagine a crude causal view, that aside from demanding that known proposition are probable enough on the evidence to satisfy some threshold, it also says that you can't know propositions that your evidence doesn't have the right kind of causal relation to. On this view you don't know that coin 1000 will not be flipped even though the threshold requirement is satisfied (assuming $k < 1 - \frac{1}{2^{1000}}$).

Moreover, below probability 1 knowledge (and justification) thresholds are famously used as premises in skeptical arguments. Such is the case in some versions of Cartesian (or closure) type arguments for which, e.g., contextualism is proposed as a remedy (Cohen 1988, is a notable example regarding lottery puzzles).

The point is, below probability 1 knowledge thresholds are rarely proposed as anti-skepticism solutions. As stated in the main text, they are often accepted implicitly by acceptance of knowledge entailing belief and the Lockean theses. See Lewis (1996), though, for minority non-skeptical view that knowledge isn't fallible and doesn't entail belief.

⁸ Since the chance that C_{1000} won't be flipped is less than 1, it seems plausible that k should be less than 1. One could claim, though, with Williamson (2009), that though the chance it won't be flipped is less than 1, the evidential probability that C_{1000} won't be flipped is 1 (all knowledge has evidential probability 1 on the $K = E$ thesis). So we do not want to deny the possibility of a probability 1 threshold view along Williamsonian lines that falsifies (5) and (6). Nevertheless, Bacon (2014) shows that in the current setting this would be a problematical move (and see footnote 16 below). Fortunately, a much less radical view is required here, one that is shared by many theorists. Aside from the reasons provided in the main text, fallibilists, will mostly agree that knowledge has a non-maximal threshold since they agree that to know a proposition doesn't require evidence that entails it. See Smith (2016) for the claim that this general "Risk Minimisation" view is prominent and widespread (a view he argues against).

⁹ As there must be in lottery cases, if we do not know that our tickets are losers before the lottery.

that C_n will not be flipped is the different probabilities, so if C_{n-1} fails to be knowledge (which by assumption it does) it must be because it is not sufficiently supported by the probabilistic evidence, and, thus, its probability is below the (contextual) threshold k .

Note that the setup makes the following two statements equivalent: C_n will not be flipped and C_{n-1} will not be flipped $\vee C_{n-1}$ will land heads. The certainty of the setup—assumption (b)—combined with the principle of knowledge equivalence—that if you know two propositions are *a priori* equivalent and you know one, you are in a position to know the other—allows us to assume that you are in a position to know the following disjunction:¹⁰

(8) C_{n-1} will not be flipped $\vee C_{n-1}$ will land heads.

Note that (7) and (8) both have a probability of $1 - \frac{1}{2^{n-1}}$ on the setup. By assumption this is greater than k (or equal to it). Note, in contrast, that (8)'s disjuncts each have a probability that is strictly lower than k . So on the k -THRESHOLD assumption, you do not know (9) and certainly not (10):

(9) C_{n-1} will not be flipped ($1 - \frac{1}{2^{n-2}}$ probability)

(10) C_{n-1} will land heads ($\frac{1}{2^{n-1}}$ probability).

Furthermore, (8) is equivalent to (and hence has the same probability as)

(11) C_{n-1} will not land tails ($1 - \frac{1}{2^{n-1}}$ probability)

and so both (8) and (11) are equivalent to (7)—that C_n will not be flipped (because if C_{n-1} doesn't land tails it follows C_n will not be flipped and if C_n won't be flipped, that can only be because a previous coin landed heads).

Note, too, that (11) has a much greater probability than that of (10). That C_{n-1} will land heads is *unlikely* because it is unlikely to be flipped at all, while that C_{n-1} will not land tails is *likely* because it is unlikely that C_{n-1} will be flipped. But although they are dramatically different in the present setting, they seem to be saying the same thing, for in a normal context that a fair coin will not land tails *implies that it will land heads*. FAIR COINS seems valid, then, because it appears to say that you can't know any particular coin will land heads. Yet it can't be taken as a statement about particular coins landing heads, since knowing a particular coin won't land tails doesn't entail it will land heads (as made vivid in DGH's scenario). To know that a coin won't land tails in this scenario is merely to know that one of a large number of fair coins will land heads, not that a particular coin will. Specifically, the proposition that *coin $n-1$ won't land tails* doesn't entail that *coin $n-1$ will land heads*. It merely entails *there will not be a sequence of $n-1$ coin flips that will all land tails* (or equivalently that *coin n will not be flipped at all*).

¹⁰ We can safely assume that known equivalences are equiknowable since even those who question equivalence closure—e.g., Yablo (2016)—will not question the present instance of it. But even if one would want to question this principle, given the equality of the probabilities, a reason needs to be given why (8) doesn't satisfy some knowledge condition that (7) does.

Moreover, (11) is where the backward induction stops. You know that C_n will not be flipped so you know the equivalent (11) and (8). But you do not know (8)'s disjuncts even though (9)—(8)'s first disjunct—is only very slightly less probable. Hence, since you do not know (9), its equivalent

(12) C_{n-2} will not be flipped \vee C_{n-2} will land heads $(1 - \frac{1}{2^{n-2}}$ probability)

is also unknowable (before the coins are actually flipped). Both (12) and (9) have the same probability that comes close, but importantly falls short of k .

This concludes our explanation for how k -THRESHOLD can be used as a reasoned rejection of (5), (6), and FAIR COINS. It boils down to a case where a disjunction's first disjunct (C_{n-1} will not be flipped) isn't quite probable enough $(= 1 - \frac{1}{2^{n-1}})$ to satisfy threshold k while with the added minute probability of its second disjunct $(= \frac{1}{2^{n-1}})$, the disjunction— C_{n-1} will not be flipped \vee C_{n-1} will land heads—satisfies k . That a disjunction's probability can satisfy a threshold while none of its disjuncts do, is the straightforward implication of the threshold view that we claimed implies the rejection of FAIR COINS.

We've also claimed that the reason FAIR COINS is intuitively compelling, is that it masks the fact that under the assumptions of the setting, saying that a certain coin won't land tails is the same as saying that there won't be a *succession of tails tosses* that will allow that coin to be tossed in the first place if (against great odds) it will be tossed, it will land heads.

Let's now inspect DGH's claim that their argument regarding knowledge can be extended to justification.¹¹ The principle they target is a variant of FAIR COINS:

JUSTIFIED FAIR COINS: If you have justification to believe that a coin is fair, and you lack justification to believe that it won't be flipped, then you lack justification to believe that it won't land tails.

By an argument analogous to that of [FAIR COINS], holding on to JUSTIFIED FAIR COINS requires denying that, in the coin example, you have justification to believe that not all of the coins will be flipped.

(DGH, p. 280)

We can be brief here since even if (rational) belief has a different (and plausibly lower¹²) threshold r , we can repeat the same (5)–(12) argument using the same (probabilistic) equivalences and differences as before. In fact, with regard to justification, we are in agreement with DGH:

Those who identify belief with confidence above a certain threshold have an independent reason to reject JUSTIFIED FAIR COINS: since the proposition that a coin won't land tails is logically weaker than the proposition that it won't be flipped, one might have justification to invest above-threshold confidence in

¹¹ We use "justified" and "rational" interchangeably. See Cohen (2016) for an argument that these should really be viewed as the same.

¹² See Hawthorne et al. (2016) for arguments that rational belief requires lower standards of evidence than knowledge does.

the former without having justification to invest above-threshold confidence in the latter.

(DGH, pp. 280–281, n. 8)

Here DGH seem to make the same point that we do above. Threshold views give independent reason to reject FAIR COINS. However, they limit the point to a very specific case: the Lockean view where belief (or perhaps justified belief) is *identified* with (justified) confidence above a certain threshold. Our argument above aims to show that a *much wider* class of views gives one independent reason to reject FAIR COINS. In particular, our argument above can be adapted to show that holding any view in which (justified) belief *requires* (justified) confidence above a certain threshold gives one good reason to reject FAIR COINS.

Moreover, any such view that puts a threshold requirement on justified belief will likely also amount to an endorsement of *k*-THRESHOLD. For if knowledge entails rational belief, as most theorists accept (not Lewis 1996, though, see footnote 6), then a threshold requirement on rational belief entails a threshold requirement on knowledge. Thus a wide range of accounts of justification will give us an independent reason to reject both JUSTIFIED FAIR COINS and FAIR COINS.

One might wonder how much our argument above depends on the sharpness of the threshold. What if there is—even in a context—no precise threshold *k* that is the minimum evidential requirement for knowledge. What if knowledge itself is vague in some cases where one’s evidence is at the threshold?

How to model such vagueness itself represents one of the more intractable philosophical questions. That *some* way of modeling a vague threshold requirement might allow us to maintain FAIR COINS is a possibility we cannot ignore. However, we want to note here that standard ways of modeling a vague threshold will not be compatible with FAIR COINS. As an example, take a supervaluationist view of a vague threshold (Fine 1975). For each precise threshold some instance of the conditional in FAIR COINS will be false. Thus, the universally quantified statement (3) will be false. So supervaluationist threshold views will also reject FAIR COINS.

3 Disjunctive threshold propositions

The disjunction we appealed to, (8), resembles what Hawthorne (2003) calls “junk disjunctive knowledge.” One can know that a disjunction $P \vee Q$ is true, he argues, even though “one of [the] disjuncts is such that if one acquires a belief in its negation, one will (or at least ought to) simply throw out the disjunction” (Hawthorne 2003, p. 72).¹³ Such are cases where one’s belief in the disjunction is solely (and rationally) based on one of its disjuncts.

¹³ The “junk disjunction” knowledge account is an expansion of Sorensen (1988) “junk conditional” account that is aimed at resolving Kripke (2011) Dogmatism puzzle. Sorensen’s account relies on Jackson’s 1979 view that introduces the technical term of “robustness” with respect to some information. In particular, Sorensen relies on the non-robustness of a conditional with respect to a conditional’s antecedent (when the information is taken to support its truth while it is believed on the basis of the antecedent’s falsity).

The present case is similar because your knowledge of the disjunction is *almost entirely* based on your belief in the truth of its first disjunct. Without the added (minute) probability that C_{n-1} will land heads, your belief in the truth of the disjunction would not amount to knowledge. Nevertheless, if you ever learned that there is a significant chance that despite the odds C_{n-1} will in fact be flipped, you wouldn't thereby be in a position to conclude by modus tollens that it will indeed land heads. You should "throw it out." So your knowledge that the disjunction is true that serves as a counterexample to FAIR COINS, shares this feature with "junk knowledge" disjunctions.

Though the usefulness of such disjunctions in deductive reasoning isn't the major issue of this paper, an analogy between the following case and FAIR COINS is instructive we think. The k -THRESHOLD condition suggest a possibility that might seem dubious. Consider the following proposition:

(13) C_{n-1} will not be flipped \vee C_{n-1} will land tails $(1 - \frac{1}{2^{n-1}}$ probability)¹⁴

Assuming that everything else remains fixed and that (13) is true due to the truth of its first disjunct (if the second disjunct is true then (11), (8) and (7) are false and hence not known, contrary to our assumption above), it may seem that you could know it. It is as probable as the disjunction (8) we considered earlier, i.e., $\text{pr}(C_{n-1}$ will not be flipped \vee C_{n-1} will land heads) = $1 - \text{pr}(n - 1$ consecutive tails tosses) = $1 - \text{pr}(n - 2$ tosses of tails and C_{n-1} landing heads) = $\text{pr}(13)$, hence it is true and satisfies k -THRESHOLD. It looks, then, like there is no obstacle in knowing this disjunction.

Yet in fact, knowing that (13) is true, isn't intuitively so smooth. We've assumed that you know that C_n will not be flipped and that you do not know that C_{n-1} will not be flipped. But if you knew both (13) and (8), you would both know that C_{n-1} will not land tails and know that C_{n-1} will not land heads (i.e., if flipped, it will land tails). So *intuitively* it doesn't seem like you can know (13) while still knowing (8) without also knowing, contra our assumption, that C_{n-1} will not be flipped at all.

Despite this formidable intuition obstacles, nothing forces us to say that you can't have this knowledge. We are, however, forced to say that if both (13) and (8) are known, then the principle of multi-premise closure is not valid. This is because assuming that you know—(8)&(13)—by conjunction introduction, (9) would be knowable by equivalence-closure. And we assumed, contrary to this, that you cannot know that C_{n-1} will not be flipped because the first coin flip sequence you know will not take place is $n - 1$ coin flips long (so you know coin n will not be flipped).¹⁵

¹⁴ Consider the case where $n = 3$. Knowing that C_n will not be flipped means that $1 - \frac{1}{2^2} = \frac{3}{4} \geq k > \frac{1}{2}$. The probability of C_2 will not be flipped or C_2 will land tails satisfies k , since it equals 1 minus the probability that C_1 will land tails and that C_2 will land heads, i.e., $\frac{3}{4}$. Of course you couldn't know both (13) and (8) when $n = 3$, unless C_1 landed heads (because then either (13) or (8) would be false) and the probability of this happening is $\frac{1}{2}$. The point we are going to make, though, is that if you know the truth of these propositions, you could not know their conjunction (unless $k = \frac{1}{2}$ in which case we would not have a proof).

¹⁵ Also, you can't know that there will be a $n - 3$ sequence of tails tosses and then two heads tosses directly, because by the conditions of the case if C_{n-2} lands heads, C_{n-1} , C_n , ..., C_{1000} , won't be flipped at all.

Moreover, by repeating the same steps (with the second disjunct replacing “heads” with “tails” in (8)), we would know that $n - 2$ coins won’t be flipped, and then that $n - 3$ coins won’t be flipped, etc. and the backward deductions that lead to an unavoidable contradiction would be reinstated. But there is no such issue if multi-premise closure fails, this kind of knowledge would be similar to having justification for believing that each lottery ticket will lose but not having justification that they will all lose.

Our commitment here to the claim that if (13) is known, then multi-premise closure is invalid, is not at all surprising. No one should be surprised that multi-premise closure and below probability 1 knowledge is a difficult combination to defend. Moreover, equally unsurprising is that the explanation for why we deny multi-premise closure in this setting is not at all an attempt to avoid skepticism. Rather, our denial of multi-premise closure goes back to why we are tempted to say that (13) could be known in the first place, that is, due to k -THRESHOLD: Each of the two propositions satisfies k but their conjunction does not. The point is that the basis for multi-premise closure denial is the same source of our denial of FAIR COINS we argued for here. In accepting that there is some threshold that allows you to know a disjunction without knowing either of its disjuncts, we have essentially accepted that you can know two conjuncts without knowing the truth of their conjunction.¹⁶

If equivalence-closure holds, we know that if there is a first coin you know won’t be flipped, then you can know the coin that proceeds it will either remain unflipped, or will land heads. And this means that you can know that this coin won’t land tails. So to claim that FAIR COINS holds would be to claim that rational belief does not have a threshold or that knowledge doesn’t entail rational belief, that equivalent propositions do not have the same probability and are not equi-knowable (equivalence-closure fails), and that knowledge itself doesn’t have a threshold. There is also the option of skepticism, of course. FAIR COINS, however, only apparently makes this option more credible.

Acknowledgements Many thanks to Matti Eklund, Anandi Hattiangadi, Harvey Lederman, and an anonymous referee for this journal for their invaluable comments. This paper was presented at the Department of Philosophy, Uppsala University, Sweden. Spectre’s research was supported by the Israel Science Foundation (Grant No. 463/12). Rothschild’s research was supported by the UK Arts and Humanities Research Council (Grant Nos. AH/M009602/1 and AH/ N001877/1).

References

- Bacon, A. (2014). Give your knowledge half a chance. *Philosophical Studies*, 2, 1–25.
- Cohen, S. (1987). Knowledge, context, and social standards. *Synthese*, 73(1), 3–26.
- Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives*, 2, 91–123, 1988. ISSN 15208583, 17582245. <http://www.jstor.org/stable/2214070>.
- Cohen, S. (2016). Theorizing about the epistemic. *Inquiry*, 59(7–8), 839–857. doi:10.1080/0020174X.2016.1208903.
- Dorr, C., Goodman, J., & Hawthorne, J. (2014). Knowing against the odds. *Philosophical Studies*, 170(2), 277–287.

¹⁶ This, then, is another reason to resist a probability 1 threshold view that accepts multi-premise closure (and divorces chance and probability). We will not pursue this line of argument here, though.

- Dretske, F. (1981). The pragmatic dimension of knowledge. *Philosophical Studies*, 40(3), 363–378.
- Fine, K. (1975). Vagueness, truth and logic. *Synthese*, 30, 265–300.
- Hawthorne, J. (2003). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Hawthorne, J., Rothschild, D., & Spectre, L. (2016). Belief is weak. *Philosophical Studies*, 173(5), 1393–1404.
- Jackson, F. (1979). On assertion and indicative conditionals. *Philosophical Review*, 88(4), 565–589.
- Kripke, S. (2011). Two paradoxes of knowledge. *Philosophical troubles: Collected papers Volume 1* (pp. 27–51). Oxford: Oxford University Press.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(1), 339–359.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Rothschild, D., & Spectre, L. (2016). A puzzle about knowing conditionals. *Noûs*. doi:[10.1111/nous.12183](https://doi.org/10.1111/nous.12183).
- Smith, M. (2016). *Between probability and certainty: What justifies belief*. Oxford: Oxford University Press.
- Sorensen, R. (1988). Dogmatism, junk knowledge, and conditionals. *Philosophical Quarterly*, 38, 433–454.
- Unger, P. K. (1975). *Ignorance: A case for scepticism*. Oxford: Oxford University Press on Demand.
- Williamson, T. (2009). Replies to critics. In D. Pritchard & P. Greenough (Eds.), *Williamson on knowledge* (pp. 279–384). Oxford: Oxford University Press.
- Yablo, S. (2016). Open knowledge and changing the subject. *Philosophical Studies*, 1–25, doi:[10.1007/s11098-016-0724-1](https://doi.org/10.1007/s11098-016-0724-1).