

KNOWING AGAINST THE ODDS

Cian Dorr, Jeremy Goodman, and John Hawthorne

§1

Here is a compelling principle concerning our knowledge of coin flips:

FAIR COINS: 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.

The idea is that the only way to be in a position to know that a fair coin won't land a certain way is to be in a position to know that it won't be flipped at all.¹

One class of putative counterexamples to FAIR COINS which we want to set aside involves knowledge delivered by oracles, clairvoyance, and so forth. A second, and more interesting, class of counterexamples involves knowledge under unusual modes of presentation. For example, if you introduce the name 'Headsy' for the first fair coin that will be flipped but will never land tails, then Headsy is arguably a counterexample to FAIR COINS: you know Headsy is fair, you know it will be flipped, and you know it will not land tails. Cheesy modes of presentation pose a challenge to a wide range of intuitive epistemological principles about objective chance (see Hawthorne and Laasonen-Aarnio 2009, §3): nevertheless, let us set them aside for the remainder of this paper.²

¹ We will treat 'For all you know, Φ ' as equivalent to 'You are not in a position to know not- Φ '. If we instead treated 'For all you know, Φ ' as equivalent to 'You don't know that not- Φ ' or 'What you know is consistent with Φ ', we would have to consider putative counterexamples to FAIR COINS in which you know that a coin will not land tails but fail to know that it will not be flipped simply because you have failed to consider whether it will be flipped.

² It is important to distinguish FAIR COINS from the stronger principle that when a coin is *in fact* fair and will be flipped for all you know, it will land tails for all you know. This principle is clearly false: suppose that you're not sure whether a coin is fair or double-headed, but know it will be flipped only if it is double-headed.

§2

Here is a case that makes trouble for FAIR COINS. 1000 fair coins are laid out one after another: $C_1, C_2, \dots, C_{1000}$. A coin flipper will begin 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. It follows that there is a smallest number n such that you know that C_n will not be flipped.³ But then C_{n-1} is a counterexample to FAIR COINS. You know that it is fair (under the ‘non-cheesy’ guise of a description specifying its position in the sequence). You know that it won’t land tails, since if it did land tails, C_n would be flipped—which you know won’t happen. But for all you know, C_{n-1} will be flipped, since n is the smallest number such that you know that C_n will not be flipped.

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.
- (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 inferences by universal instantiation and modus ponens. But (4) is

³ We do not claim that there is any n such that you *definitely* know that C_n will not be flipped and *definitely* fail to know that C_{n-1} will not be flipped: this stronger claim is implausible given the vagueness of ‘know’.

obviously true, since C_1 will be flipped. (2) is hard to deny, given that you know the setup.⁴ 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. Since you also know that this coin is fair, we have a counterexample to FAIR COINS.⁵

§3

The puzzle deepens. The following principle is even more compelling:

BIASED COINS: If you know that a coin is heavily biased towards tails, and for all you know it is going to be flipped, then for all you know it will land tails.

No matter how extreme the bias, there could be a sequence of biased coins long enough that we know that they won't all be flipped (keeping the setup as before). The above argument will then generate a counterexample to BIASED COINS.

§4

Your first instinct may be to give up the anti-skeptical (1) in order to save FAIR/BIASED COINS. The trouble is that it is hard to prevent a skeptical line on long sequences of tails from extending to most of what we take to be mundane

⁴ Since the proposition that C_{n-1} will not land tails is an obvious logical consequence of the proposition that C_n will not be flipped together with the setup, one could defend (2) by appealing to the principle that what we are in a position to know is closed under obvious logical consequence. However, even if one does not accept that closure principle in full generality, there is little promise in the view that while it is possible for a person to know the setup and that C_{1000} won't be flipped, it is impossible for a person to know this while satisfying (2).

⁵ In deriving the negation of (3) from (1), (2) and (4) we are using *reductio ad absurdum*, a rule which some philosophers regard as invalid in the presence of vagueness. But even these philosophers should, we think, regard the above argument as constituting a good reason to *reject* (3), and with it FAIR COINS, even if it is not a good reason to accept its negation.

knowledge of the future. Consider for example a particular leaf on a maple tree that sheds all of its leaves every winter. If you know anything at all about the future, you know that come February the leaf will no longer be on the tree. Now divide the time between now and February into hour-long intervals. Suppose you know that for each hour, if the leaf is still on the tree at the beginning of that hour, the chance at the beginning of the hour that it will still be on the tree at the end of the hour is high. Presumably, anyone who accepts BIASED COINS will also accept the following principle:

AUTUMN LEAF: For all hour-long intervals h : If you know that leaf-shedding is a chancy process of the kind just described, and for all you know the leaf will still be on the tree at the beginning of h , then for all you know the leaf will still be on the tree at the end of h .

The analogy between AUTUMN LEAF and BIASED COINS should be clear. But AUTUMN LEAF is incompatible with the anti-skeptical assumption that you can know both that the leaf won't be on the tree in February and that leaf-shedding is a chancy process of the relevant kind.⁶ For the same reasons as in the coin case, there must be a first hour such that you are now in a position to know that the leaf will not be on the tree at the end of it. Since it is the first such hour, you are not in a position to know that the leaf will not be on the tree at the beginning of it. And since you know that if the leaf hasn't fallen by the beginning of that hour it will be objectively improbable that it falls by the end of that hour, we have a counterexample to AUTUMN LEAF.

So denying (1) has sceptical ramifications far beyond cases involving coin-flips and similarly artificial chancy processes. Indeed, the skeptical solution to our puzzle threatens to engulf a fair amount of our knowledge of the past and present

⁶ We are not suggesting that it is inevitable that the mechanisms of leaf-shedding do in fact work like this, though it is not completely unrealistic. In §6, we suggest that a puzzle can be generated by a weaker premise, namely that *for all you know*, the chance at the beginning of each hour that the leaf will stay on the tree for another hour is high.

along with most of our knowledge of the future. For example, if you can't now know that the leaf won't still be on the tree in February, then presumably you won't be able to know it in February either, if you haven't seen or heard about the leaf in the meantime.⁷

§5

Some philosophers have already convinced themselves that we know hardly anything about the future. However, many such philosophers think that we nevertheless have many *justified beliefs* about the future: for example, a justified belief to the effect that a certain leaf will fall before February. While those who hold this combination of views will be unfazed by the skeptical consequences of principles like FAIR COINS and AUTUMN LEAF, they will still have to reject formally analogous, and similarly attractive, principles formulated in terms of justified belief, such as:

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 §2, holding on to JUSTIFIED FAIR COINS requires denying that, in the coin example, you have justification to believe that not

⁷ Some contextualists may be tempted to opt for a familiarly attenuated skepticism according to which, in certain contexts in which we are thinking statistically or in which error possibilities are salient to us, 'know' takes on a meaning under which the anti-skeptical premise (1) is false. But such contextualists standardly think that there are plenty of contexts in which the anti-skeptical premise is true. Since contexts in which (1) is true will typically be ones in which FAIR COINS is false (since they will typically be ones in which (2) and (4) are true), this sort of contextualism does not threaten the interest of our argument.

all of the coins will be flipped. And by an argument analogous to that of §4, this lack of justification will extend to many mundane beliefs about the future.⁸

§6

How broadly does the failure of FAIR COINS ramify? Let us survey some principles concerning knowledge and objective chance to see which are consistent with the falsity of FAIR COINS and which are not.

Here is one attractive principle that is perfectly consistent with the falsity of FAIR COINS, and indeed with everything we have said so far:

KNOWN UNLIKELYHOOD: If you know that there is a substantial objective chance that P, then for all you know, P.

Note that in the coin-flipping example, your knowledge that C_n won't land tails is no counterexample to KNOWN UNLIKELYHOOD, since C_n has a very low chance of being flipped at all, and hence an even lower chance of landing tails.

There are several natural strengthenings of KNOWN UNLIKELYHOOD that are still consistent with the falsity of FAIR COINS:

ACTUAL UNLIKELYHOOD: If there is a substantial objective chance that P, then for all you know, P.

KNOWN FUTURE UNLIKELYHOOD: If you know that there is or will be a substantial objective chance that P, then for all you know, P.

⁸ 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 the former without having justification to invest above-threshold confidence in the latter.

ACTUAL FUTURE UNLIKELIHOOD: If there is or will be a substantial objective chance that P, then for all you know, P.

Our anti-skeptical judgments about the coin case are thus compatible with a range of attractive views according to which facts about objective chances place stringent constraints on what we can know about the future.⁹

On the other hand, the following natural strengthening of KNOWN UNLIKELIHOOD does entail FAIR COINS, and must therefore be rejected:

KNOWN CONDITIONAL UNLIKELIHOOD: If you know that there is a substantial objective chance that P conditional on Q, and for all you know, Q, then for all you know, P.

The route from KNOWN CONDITIONAL UNLIKELIHOOD to FAIR COINS is straightforward. If you know that a coin is fair, then you are in a position to know that the objective chance that it will land tails conditional on it being flipped is 50%. Since 50% is certainly a 'substantial' chance, KNOWN CONDITIONAL UNLIKELIHOOD entails that if for all you know the coin will be flipped, then for all you know it will land tails.

KNOWN CONDITIONAL UNLIKELIHOOD is, however, subject to counterexamples of a more straightforward character than those which refute FAIR COINS. Suppose Jack is acquainted with tens of thousands of people, one of whom, Jill, will be struck by lightning during the coming week. Let P be the proposition that Jack won't be struck by lightning in the coming week. Given anti-skepticism about the future, we may assume that Jack knows P. Let Q be the proposition that either Jack or Jill will be struck by lightning in the coming week. Since Q is true, it is true for all Jack knows. But Jack knows that there is a substantial objective chance of P conditional on Q, since he knows that he and Jill have a similar chance of being

⁹ For an important line of argument against KNOWN UNLIKELIHOOD, based on the principle that knowledge can be extended by deduction from known premises, see Hawthorne 2004, §4.6 and Williamson 2009.

struck by lightning in the coming week. Since this situation is inconsistent with KNOWN CONDITIONAL UNLIKELIHOOD, that principle leads immediately to rampant scepticism about the future.

Here is a different natural strengthening of KNOWN UNLIKELIHOOD which also yields FAIR COINS, but for which the case of Jack and Jill is not a counterexample:

POSSIBLE FUTURE UNLIKELIHOOD: If for all you know, there is or will be a substantial objective chance that P, then for all you know, P.¹⁰

We can derive FAIR COINS from POSSIBLE FUTURE UNLIKELIHOOD as follows. Suppose that there is a coin that you know is fair and that for all you know will be flipped. Then for all you know, there will be a time, namely when it is flipped, at which the objective chance that it will land tails is 50%, which is substantial. So by POSSIBLE FUTURE UNLIKELIHOOD, the coin will land tails for all you know.

POSSIBLE FUTURE UNLIKELIHOOD entails not just FAIR COINS but the following stronger claim, which applies not only to coins that you know to be fair but also to coins that are fair for all you know:

STRONG FAIR COINS: If for all you know, a coin is fair and going to be flipped, then for all you know, it will land tails.

For similar reasons, POSSIBLE FUTURE UNLIKELIHOOD entails the following strengthening of AUTUMN LEAF:

¹⁰ Note that if POSSIBLE FUTURE UNLIKELIHOOD is not to have the absurd consequence that only people with the concept of objective chance can know anything, we had better understand its antecedent in such a way that it is not automatically true of people who lack the concept of objective chance. Such understandings of 'for all you know' are not unfamiliar—surely, even if you lack the concept of a hexagon, there is a natural reading of 'For all you know, your hand is a hexagon' on which it is false. Similar points apply to 'in a position to know'.

STRONG AUTUMN LEAF: If a future hour-long interval h is such that for all you know, there is a substantial objective chance at the beginning of h that a certain leaf will not have fallen by the end of h , then for all you know, that leaf will not have fallen by the end of h .

Claims like STRONG AUTUMN LEAF will more obviously lead to widespread skepticism about the future than principles like AUTUMN LEAF: the latter principles rule out knowledge about the future only in cases where we know certain facts about the underlying chancy mechanisms, and it is not obvious how easy it is to acquire such knowledge. It would therefore be interesting if there were some simple general principle (other than the problematic KNOWN CONDITIONAL UNLIKELIHOOD) which entailed FAIR COINS and AUTUMN LEAF without entailing the above strengthened versions. But we haven't found any such principle. If there is no such principle, the skeptical cost of maintaining FAIR COINS in a principled way will be very high indeed.

§7

Even those who don't mind the idea that we know very little about the future should be wary of POSSIBLE FUTURE UNLIKELIHOOD. For even if we don't know much about what will in fact happen, we seem to know quite a lot about the objective chances of different possible happenings. (Indeed, those who claim that we know little about the future often try to make that claim easier to swallow by maintaining that in cases where we are initially tempted to say that we know that something will happen, we really do know that it has a very high objective chance of happening.) But if POSSIBLE FUTURE UNLIKELIHOOD is true, it is hard to see how we could ever learn anything non-trivial about objective chances.

Surely, if you could ever learn anything non-trivial about objective chances, you could learn that a certain double-headed coin is not fair by flipping it repeatedly, seeing it land heads each time, and eventually inferring that it is not fair. In any such case, there must be a first flip of the coin after which you are in a

position to know that the coin is not fair. Before that flip, the coin was fair for all you knew. So for all you knew, the proposition *that the coin was fair and about to land heads* had a substantial (viz. 50%) chance of being true. Given POSSIBLE FUTURE UNLIKELIHOOD, it follows that for all you knew then, the coin was fair and about to land heads. And yet, upon seeing it land heads, you were somehow able to infer, and thereby come to know, that it was not fair! This is quite odd, and in conflict with the following intuitive principle about inferential knowledge

INFERENCE ANTI-DOGMATISM: If for all you know, P and not-Q, then you cannot come to know Q just by learning P and inferring Q from P and things you already knew.

So, given INFERENCE ANTI-DOGMATISM, POSSIBLE FUTURE UNLIKELIHOOD will lead not only to a wide-ranging scepticism about the future, but to scepticism about current objective chances.¹¹

We don't want to overstate the significance of this argument. Although we find INFERENCE ANTI-DOGMATISM quite plausible, many of those who think that we know little about the future already have reason to reject it. We have in mind packages according to which, although we know nothing about the future whose objective chance is less than one, we can know a fair amount about the past and present by making ordinary non-deductive inferences.¹² For example you can, on

¹¹ Note that INFERENCE ANTI-DOGMATISM does not rule out the kind of 'perceptual dogmatism' according to which you can come to know that a certain surface is red by looking at it even if, for all you knew before looking, the surface was white but illuminated in such a way as to look red. (A justification-theoretic analogue of this thesis is defended by Pryor (2000) and attacked by White (2006).) INFERENCE ANTI-DOGMATISM merely says that in such a case, you cannot come to know that the surface is red by means of an *inference* from the fact that it looks red.

¹² Although POSSIBLE FUTURE UNLIKELIHOOD does not rule out all knowledge of propositions whose objective chance is less than one, we suspect that friends of POSSIBLE FUTURE UNLIKELIHOOD will find it hard to stake out a principled view about when propositions whose chance is less than one can be known. For example, although POSSIBLE FUTURE UNLIKELIHOOD rules out knowledge that a sequence

the basis of a digital thermometer's reading 45 degrees, infer and thereby come to know that the temperature is between 40 and 50 degrees. And you have this knowledge despite the fact that, prior to the measurement, there was a tiny but nonzero objective chance that the temperature would suddenly fluctuate up to 55 degrees while, owing to some equally improbable compensating fluctuation inside the thermometer, it would nevertheless read of 45 degrees. But because of this nonzero chance, it was true for all you knew before the measurement that the thermometer would read 45 degrees while the temperature was not between 40 and 50 degrees. This package of commitments clearly requires rejecting INFERENTIAL ANTI-DOGMATISM. In general, INFERENTIAL ANTI-DOGMATISM has a tendency to force those who endorse limited forms of skepticism about knowledge of the future into more wide-ranging forms of skepticism about inferential knowledge of the present and past, including knowledge of the objective chances.

§8

Let's return to our original coin-flipping case. Notice that every coin that will in fact be flipped is a coin that will come up tails for all you know. So the case is no counterexample to the following principle:

WEAK FAIR COINS: If a coin is fair and will be flipped, then for all you know it will come up tails.

We find this principle extremely plausible. Are there any compelling arguments against it?

of *successive* flips of fair coins will not all come up tails, it does not rule out our knowing in some cases that a collection of fair coins that will be flipped *simultaneously* will not all come up tails. But it seems bizarre to suppose that the difference between a protocol in which coins are flipped successively and one in which they are flipped simultaneously could actually have this kind of epistemic significance.

Let's assume that, if WEAK FAIR COINS is true, then it is something that you can know to be true, and moreover, know to be true while in the coin case and knowing every relevant fact that you are in a position to know. In that case, it is plausible that each coin that you *know* you know won't land tails is a coin you know won't be flipped, since the fact that it won't be flipped is an obvious consequence of the known facts that WEAK FAIR COINS is true and that you know the coin won't land tails.¹³ So, if WEAK FAIR COINS is true, then plausibly so is

(5) For each coin C_n : If you know that you know that C_n will not land tails, then you know that C_n will not be flipped.

(5) is like the problematic principle (3) from our original inconsistent tetrad, except that its antecedent contains two iterations of knowledge rather than one. (3) thus follows from (5) together with the following principle:

(6) For each coin C_n : If you know that C_n will not land tails, then you know that you know that C_n will not land tails.

Since we reject (3), we must therefore reject either WEAK FAIR COINS or (6).

Luckily for WEAK FAIR COINS, there are strong grounds for rejecting (6). Independently of any general theoretical commitments, in imagining the example it strikes us as implausible to suppose that the first coin that you know won't land tails is one that you know you know won't land tails. Moreover, whatever temptation there might be to accept (6) seems to derive from its being an instance of the KK principle, according to which everything you know is something you know you know. Since the KK principle is widely discredited (see Williamson 2000,

¹³ One could resist this argument by denying that your knowledge in this case would be closed under obvious logical consequence. But it is hard to think of a principled reason why the factors responsible for failures of closure would have to present whenever you found yourself in the coin example while knowing WEAK FAIR COINS.

chapter 5), we see no compelling reason to accept (6), and thus no compelling reason to reject WEAK FAIR COINS.

Here is a second argument against WEAK FAIR COINS. You might think that your ability to know that there won't be 1000 tails in a row is explained simply by the fact that it is (a) true and (b) known to have very high objective chance. But if this combination is sufficient for knowledge, then it will be possible to know that C_{1000} won't land tails even in a world where it will be flipped and land heads (after all 999 preceding coins land tails). In such a world, C_{1000} would therefore be a counterexample to WEAK FAIR COINS.¹⁴

However, we think there are strong grounds to deny that the combination of truth with known high chance suffices for knowledge in these sorts of cases. For one thing, it simply strikes us as implausible that in a world where C_{1000} is tossed and lands heads, we can know in advance that it won't land tails. This judgment is reinforced by the plausibility of the following 'margin for error' principle: if a sequence of possible outcomes of flips of fair coins differs at only one position from the sequence of the actual outcomes of the flips of those coins, then for all you know

¹⁴ The proposal that the combination of true belief with known high chance suffices for knowledge can be weakened in two ways without disrupting the argument against WEAK FAIR COINS. First, one might add a further requirement that the propositions in question are ones whose negations would be *remarkable* (in a sense of 'remarkable' on which not just any low-chance truth is remarkable: for example, it would not be remarkable for the outcomes of a series of ten coin-flips to be HTTHTHHTTTH). Second, one might restrict the generalization to cases where the complete truth about the relevant subject matter is not itself remarkable: if a certain coin will in fact land heads a hundred times in a row, it is far from clear that one could know that it won't land tails a hundred times in a row. Even so weakened, this sufficient condition for knowledge can still be used to argue against WEAK FAIR COINS. Whatever 'remarkable' means, there will be a least n such that it would be remarkable, in our coin-flipping case, for the n th coin to be flipped and land tails. Suppose that in fact the n th coin is flipped but lands heads, so that the truth about the coins is unremarkable. Then, according to the weakened sufficient condition for knowledge, you can know that the n th coin won't land tails, since this proposition is true, is known to have a very high chance, and has a remarkable negation.

that sequence will obtain (cf. Williamson 2000, chapter 5). We therefore reject this second argument against WEAK FAIR COINS.¹⁵

§9

So although we reject FAIR COINS, we accept WEAK FAIR COINS. More generally, while we reject POSSIBLE FUTURE UNLIKELIHOOD (from which FAIR COINS can be derived), we are still attracted to ACTUAL FUTURE UNLIKELIHOOD (from which WEAK FAIR COINS can be derived). The relationship between these two general principles fits a familiar pattern: the latter says that a certain phenomenon (namely, present or future low objective chance) is incompatible with knowledge, while the former says that the mere epistemic possibility of that phenomenon is incompatible with knowledge. But once we reject KK, we must recognize that epistemic principles cannot always be strengthened in this way. After all, *ignorance* is incompatible with knowledge, but given the falsity of KK, the mere epistemic possibility of ignorance is not incompatible with knowledge. Similarly, even if the presence of nearby fake barns is incompatible with knowing that one is looking at a barn, it is much less plausible to insist that one cannot know that one is looking at a barn unless one is in a position to *know* that there are no fake barns nearby. And many philosophers would agree that even if certain kinds of perceptual unreliability are incompatible with perceptual knowledge, possession of such knowledge does not require that one is in a position to know that one is perceptually reliable. The perspective we recommend is one according to which POSSIBLE FUTURE UNLIKELIHOOD is an unacceptably tendentious strengthening of ACTUAL FUTURE UNLIKELIHOOD—whatever the ultimate fate of the latter, we shouldn't be put off it by the failure of the former. FAIR COINS strengthens WEAK FAIR COINS in an

¹⁵ The margin for error principle is also clearly inconsistent with the weaker sufficient condition for knowledge discussed in note 14: in any case where the truth about a long sequence of future coin-flips is just barely unremarkable, differing from a remarkable sequence at only one position, the weaker sufficient condition entails that we can know that the remarkable sequence of outcomes will not be actualized.

analogous way: our considered view is that while FAIR COINS is immediately gripping, WEAK FAIR COINS is the salient truth in the vicinity.¹⁶

¹⁶ Thanks to Andrew Bacon, Stewart Cohen, Jane Friedman, Peter Fritz, Harvey Lederman, Jeff Russell, and Timothy Williamson, and to the participants in a workshop at All Souls College and seminars in Oxford and Princeton.

References

- Greenough, P. & Pritchard, D. (Eds.) (2009). *Williamson on Knowledge*. Oxford: Oxford University Press.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, J. and Lasonen-Aarnio M. (2009). Knowledge and Objective Chance. In Greenough and Pritchard 2009, 92–108.
- Pryor, J. (2000). The Skeptic and the Dogmatist. *Noûs*, 34, 517–549.
- White, R. (2006). Problems for Dogmatism. *Philosophical Studies*, 131, 525–557.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- (2009). Reply to John Hawthorne and Maria Lasonen-Aarnio. In Greenough and Pritchard 2009, 313–329.