

Giving Your Knowledge Half a Chance

Andrew Bacon*

ABSTRACT: 1000 fair causally isolated coins will be independently flipped tomorrow morning and you know this fact. I argue that the probability, conditional on your knowledge, that any coin will land tails is almost 1 if that coin in fact lands tails, and almost 0 if it in fact lands heads. I also show that the coin flips are not probabilistically independent given your knowledge. These results are uncomfortable for those, like Timothy Williamson, who take these probabilities to play a central role in their theorizing.

Consider the following thought experiment. 1000 coins will be flipped tomorrow morning and you know this. You also know two further facts about these coins. Firstly, you know that the coins are fair – they have exactly as much objective chance of landing heads as tails. Secondly you know that they are causally isolated from one another – you know, let us suppose, that each of the coins are distributed over a number of different planets located in the Milky Way. There is no particular way of ordering the coins, but we may suppose that they have arbitrarily been assigned numbers between 1 and 1000 for convenience.

Let H_n and T_n represent the propositions that the n th coin will land heads and tails respectively. The following question will be the focus of this paper:

QUESTION: what is the probability, given your knowledge, that a given coin will land tails.

The notion of probability given your knowledge plays an important role in the theory developed by Williamson in [8] (chapter 10.) I will follow him in calling these ‘evidential probabilities’. The question, however, can be made intelligible without assuming the theoretical commitments of this particular theory.¹ In order to make sense of this notion it suffices to assume that each agent is associated with at least one, and possibly a collection, of candidate ‘ur-priors’: probability functions that represent the agents opinions prior to obtaining any evidence at all. An agent’s evidential probability function at t is then determined by conditioning their prior probability function, Pr , on the conjunction

*Many thanks to Kenny Easwaran, Jeremy Goodman and John Hawthorne for helpful discussions on earlier drafts of this paper.

¹Indeed, some aspects of Williamson’s framework are completely unnecessary for this discussion. Most notably, his assumption that there is exactly one theoretically significant ur-prior.

of propositions that that agent knows at t (if there is more than one ur-prior there may be more than one evidential probability at t .)² The notion of the probability of a proposition conditional on your knowledge, therefore, is something many people, including traditional Bayesians, can make sense of.³ The most fruitful area of dispute, to my mind, is not whether they exist, but whether they are theoretically significant – are they the credences one ought to have, do they represent the betting odds one ought to adopt, do they measure, in any intuitive sense, the degree to which a proposition is justified and so on.

Before we can address the problem in a precise way we must settle some preliminary issues that will help us articulate the problem. The first of these concerns the way in which we will represent the relevant epistemic possibilities. It is natural to think that our knowledge about the weather, politics, the solar system and so on has no bearing on the evidential probabilities of T_n and H_n for any n : learning when the Napoleonic wars took place should not make it more or less likely given my knowledge that the 121st coin will land heads.⁴ Given this assumption we can set aside many aspects of our epistemic state as irrelevant to the present puzzle. The area of logical space characterised by our set up (the number of coins, their chances and so on) can be treated as partitioned into equiprobable propositions describing the exact sequence of outcomes for each of the coins. These propositions will be equiprobable in the sense that they have equal chance; this ensures, if we assume the Principal Principle (discussed below), that they also have equal probability according to the prior Pr . Our knowledge, in this situation, can be represented as a disjunction of elements of this partition.⁵ We shall represent each member of this partition by a function σ mapping the numbers less than or equal to 1000 onto one of two possible outcomes, T and H . $\sigma(n)$ represents whether the n th coin lands heads or tails. I'll represent the sequence that will actually happen with the symbol α . One's evidential probability is the result of conditioning Pr on this area of logical space minus the sequences we know won't happen (that all the coins land heads, or

²This definition makes sense even if the agent doesn't know the conjunction of the propositions she knows. Thus, although evidential probabilities so defined are available even to anti-closure epistemologists, it may be an uninteresting notion. For example, theorists who claim that we know that that each ticket in a fair lottery won't win, bar the winning one, will assign an evidential probability of 1 to the winning ticket and 0 to the others according to the above definition. I will therefore exclude these kinds of views from consideration in what follows – such views would have little use for evidential probabilities anyway.

³Those who find the idea of a primordial ur-prior too fantastic, even when conceived as an abstraction from present credences, can make do with the notion of credence conditional on their knowledge. In order for this to be well defined one must assume (quite plausibly) that the agent is not certain that her knowledge is false. One ought to be able to reconstruct everything I say about evidential probabilities using this notion.

⁴The following observations might lead one to question this assumption, indeed I suspect myself that it is probably false. The use of this assumption does not diminish the interest of this thought experiment however, which, among other things, is intended to demonstrate that assumptions like this are not true in general to those who would initially accept the assumption.

⁵If our knowledge about politics, say, could make it more or less likely that that the first coin will land heads then these simplifications are not innocent.

whatever it may be.)⁶

The proposition representing the set-up (the number of coins, their chances and so on) can thus be represented by the set of all such sequences, the conjunction of my knowledge can be represented by a subset, K , of such sequences (the sequences that, for all I know, will represent the outcomes of the flips), T_n and H_n with a set of sequences that agree in the n th place, and so on. Given our above assumptions the evidential probability of T_n , which I'll write $Q(T_n)$, is $Pr(T_n | K)$, which can be calculated by dividing the cardinality of $T_n \cap K$ by the cardinality of K : $Q(T_n) = \frac{|T_n \cap K|}{|K|}$. I'll call a sequence epistemically possible if it belongs to K .

The main conclusions of this paper will be drawn out in three stages, accumulating premises along the way. After some initial remarks we begin in section 2 by assuming (along with the Principal Principle) that one knows that not all the coins will land heads. On this assumption it is shown that the coin flips are not probabilistically independent of one another when probabilities are understood in Williamson's way. This suggests that the results of some coins can be evidentially relevant to the outcomes of other (distinct) coins. In section 3 I defend two additional assumptions regarding what happens to your knowledge when you learn the outcomes of some of the flips. From these assumptions I show that some of the coins do not have an evidential probability of $\frac{1}{2}$ of landing heads. In section 4, I add another assumption about how knowledge updates and use this to argue that the evidential probability of a coin landing heads (or tails) is almost 1 if the coin in fact land heads (tails). The assumptions taken together entail a very natural 'margin for error' model of our knowledge in the case described, and in section 5 I compare this model with some other accounts of the lottery paradox. I conclude with a brief discussion of the hypothesis that evidence ought always have a chance of 1 of obtaining.

1 Initial Assumptions

In what follows I will make two assumptions that will significantly constrain our discussion. The first of these, while familiar, requires a few remarks.

THE PRINCIPAL PRINCIPLE: for any rational ur-prior, Pr , and evidence E that is admissible at t , $Pr(A | E \wedge Ch_t(A) = x) = x$ provided E is compatible with $Ch_t(A) = x$ (here Ch_t represents the chances at t .)

'Admissible evidence' is a term of art. According to Lewis's rough characterisation, evidence is admissible if it is 'the sort of information whose impact on credence about outcomes comes entirely by way of credence about the chances of those outcomes' (see [6] p92). Paradigm examples of evidence that is admissible

⁶If there are multiple ur-priors among which it is indeterminate which represents the agents uninformed opinions, they will all agree on this partition conditional on this area of logical space provided each candidate ur-prior satisfies the Principal Principle. Multiplicity in the representation of our initial credences therefore does not affect our conclusions about evidential probabilities.

at t include evidence that is historical and concerns only matters of particular fact that occurred prior to t , and evidence that is about what the chances would have been at a time if one of these historical propositions describing the world up to that time had obtained. Propositions directly about particular matters of fact later than t , i.e. information about the future, are archetypical examples of information that is inadmissible at t .

If K is the conjunction of my knowledge at time t , my evidential probability will be $Pr(\cdot \mid K)$ for some ur-prior Pr . By hypothesis you know that the coins are fair, so the evidential probability that the n th coin lands tails is: $Pr(T_n \mid K) = Pr(T_n \mid K \wedge Ch_t(T_n) = \frac{1}{2})$. Note, however, that we cannot immediately infer from the Principal Principle that the evidential probability that the n th coin lands tails is $\frac{1}{2}$. This is because, setting aside radical skepticism about the future, K is not admissible in Lewis's sense. I know, for example, that I'm having pasta for dinner and this is a proposition that is about the future and is therefore not admissible. What is more, since this fact has (and I know it to have) an objective chance of less than 1, my evidential probability will assign a probability of 1 to a proposition I know to have a chance of less than 1.⁷ Simply knowing the chances does not guarantee that your evidential probabilities match the chances. Thus the quick and simple route to answering the question we opened with, via the Principal Principle, is blocked.

The second assumption I shall make is:

ANTI-SKEPTICISM: There is at least one sequence of heads and tails, s , (say, the sequence of all heads) such that you know that s won't happen (or, at the very least: that s won't happen is entailed by things that you know.)

This assumption is a fairly straightforward anti-skeptical claim, and I think it is plausible in its own right. That said, it is very natural to draw analogies between the present case and the widely discussed lottery paradox. It is well known that, in conjunction with a closure principle for knowledge, ignorance about lottery propositions quickly leads to widespread skepticism (see Vogel [7] and Hawthorne [4].) Surely (the thought goes) I know I won't be living in a mansion next year even though I know that I'd buy a mansion if I won the lottery. Yet if I don't know whether I'll win the lottery and my knowledge is closed under obvious consequences, I can't know both these things.

Closure is an important component of Williamson's theory so, at least as far as our primary target is concerned, this assumption can be granted. Nonetheless, the most straightforward way to deny this assumption without resorting to a thorough going skepticism is to resist single premise closure, so it would cost us a significant amount of generality to ignore this possibility.

⁷This fact is much more general: since many of us know that we live in a quantum mechanical universe where most propositions about the future have a non-zero chance of being false, evidential probabilities will frequently come apart from known chances (see [5] Hawthorne & Lasonen.)

For what it is worth, I find the costs of these sorts of responses to skeptical arguments to be high.⁸ But rather than trying to take on the tricky business of defending single premise closure, let me instead show that the puzzle can arise without assuming it. In fact I think there are three reasons why single premise closure is irrelevant here. The first two are slightly more tentative, but certainly worth mentioning. To my ear, at least, it seems plausible that anti-skeptical considerations can directly motivate ANTI-SKEPTICISM without the detour through knowledge of ordinary propositions (i.e. propositions that are not about lotteries) and closure. By way of analogy, imagine I have selected the following ten digit password for my bank: 4538627916.⁹ In the good, non-skeptical, cases – i.e. excluding cases where, say, there are millions of hackers attempting to guess my password, or a single person guesses my password by fluke – it seems clear that I know that my password is secure. In these cases, if I find out that a hacker has attempted a single guess at my password then I will know that it failed. The intuition that I know the hacker didn't guess my password (i.e. 4538627916) is direct and not a consequence of closure and other intuitions. Moreover, this case is exactly analogous to the present one except that we have 10 independent guesses with 10 possible values rather than 1000 with 2.

A second reason to think that we can directly rule out some sequences of coins in advance is that it would otherwise be very hard to account for knowledge acquired by induction. Consider a variant scenario in which I don't know that the coins are fair, but my evidential probabilities are equally divided between the hypothesis that all the coins are fair and that they are all double headed. If I were to observe all 1000 coins land heads, it seems, I ought to be able to knowledgeably conclude that the coins are double headed. If I couldn't acquire knowledge this way I don't see how one could ever acquire knowledge about chances from induction. Yet if the proposition that all the coins are fair and will land heads is an epistemic possibility before I see all the coins land heads, it will surely also be an epistemic possibility after I see them all land heads. Thus I am not in a position to know that the coins are double headed after making these observations unless I know beforehand that the coins won't all land heads if they're fair. Of course, there are extremely rare circumstances in which the coins are fair and do all land heads. In these rare cases we wouldn't know that the coins are double headed (because they're not); but it would be a mistake to infer that in good cases, like our own, people can't acquire knowledge from induction in this way.

The final point, and I think this is conclusive, is that rejecting ANTI-SKEPTICISM is tantamount to accepting a thorough going skepticism. The puzzle can be reformulated so as not to rely on closure at all: the parenthetical portion of ANTI-SKEPTICISM – that we merely know some things that jointly entail that not all of the coins will land heads – is enough. The crucial fact to highlight in this regard is that whether someone knows a proposition, or

⁸See the discussion of closure in chapter 1 of [4].

⁹I owe this example to Kenny Easwaran.

merely knows some propositions that entail that proposition, makes no difference as far as evidential probabilities, as defined above, are concerned. Due to the properties of probability functions anything entailed by things you know has an evidential probability of 1 even if you don't know the entailed proposition. Thus, for example, if I know both that I won't be living in a mansion next year and that I'd be living in a mansion if I won the lottery, then the evidential probability that I'll win the lottery is 0. Anti-skeptical considerations, then, provide us with a strong reason to accept at least the parenthetical disjunct of ANTI-SKEPTICISM, and this is sufficient to show that there is at least one sequence that has an evidential probability of zero (and this fact is all I shall need to make the observations about evidential probabilities that follow.) Resisting this premise, then, is tantamount to accepting a fairly widespread skepticism that infects not just our judgments about lottery propositions, but a great deal of our everyday knowledge as well.

2 Probabilistic Independence

While these preliminary assumptions should not be too controversial, I shall argue that they have several surprising upshots for the notion of evidential probability. The first of these consequences, which I shall discuss here, states that the coins are not mutually probabilistically independent of one another. A collection of events are mutually probabilistically independent of one another if the probability of their conjunction is the result of multiplying the probabilities of the conjuncts.

More precisely, given our two assumptions above, one can show that either the coin flips are not probabilistically independent of one another or that you're in a position to know how some coin will land (the result follows since we clearly are not in a position to know how any coin will land.)

Proof. Let us suppose, without loss of generality, that we know that the coins won't all land heads. If the H_n propositions are jointly probabilistically independent of one another then $Q(\bigcap_{i \leq 1000} H_i) = Q(H_1) \cdot Q(H_2) \cdot \dots \cdot Q(H_{1000})$. Since the former number is 0 it follows that $Q(H_i) = 0$ for some i – which is to say that you know how the i th coin lands (or, if we are not assuming closure, we have an evidential probability of 1 concerning the outcome of one of these coins) – or that the coins are not jointly probabilistically independent.¹⁰ \square

Say that one proposition, B , is evidentially supported by another, A , if it is more probable on A than it is probable; in other words, if the conditional evidential probability of B on A is greater than the evidential probability of B . Those sympathetic to the framework of evidential probabilities will no doubt take this notion to be a useful and significant measure of how much a piece of evidence supports a hypothesis.

¹⁰It is worth noting that, for all I've said, it's still possible that the coins are pairwise independent of one another.

The above result has some surprising consequences for views of this kind. Suppose that c is a coin and that X is a set of coins that does not contain c . Intuitively it does not seem as though the outcome of c 's flip could be evidence for or against any hypothesis concerning how the coins in X land (such as the hypothesis that they land HTHHT, for example.) To make the point vivid suppose that c is a coin here on Earth and X are a set of coins on Alpha Centauri. The proposition that c will land heads surely cannot constitute evidence for or against the hypothesis that the coins on Alpha Centauri will land HTHHT. The outcome of c has positively no bearing on the outcomes of the coins on Alpha Centauri – one cannot simply improve one's epistemic position with respect to how the coins on Alpha Centauri landed by flipping a coin and observing the outcome over here.

Analogous theses seem to apply to credences. If I were to learn that a coin on Earth had landed heads I should remain just as confident as I was before about how the coins on Alpha Centauri landed. Yet one's credences ought to be regulated by evidential support: if some evidence evidentially supports a hypothesis, then one ought to become more confident in that hypothesis upon acquiring that evidence.

However the fact that the coins are not probabilistically independent of one another entails that there will be situations of exactly this kind. In particular, let n be the largest number such that $Q(\bigcap_{i \leq n} X_i) > 0$ (such an n must exist since $Q(X_1) > 0$ and $Q(\bigcap_{i \leq 1000} X_i) = 0$.) Then $Q(\bigcap_{i \leq n} X_i | X_n)$ is well defined, and identical to 0, while $Q(\bigcap_{i \leq n} X_i) > 0$.¹¹

3 Closeness

Note that the fact that the coins are not probabilistically independent of one another does not mean that the probability of T_n is anything other than a half for each n . For example, if you knew the coins would all land the same way but didn't know which way they'd all land (i.e. if $K = \{TTTTTT...T, HHHHHH....H\}$) then the probability of T_n and H_n would be a half for each n . Here the failure of probabilistic independence is particularly vivid: conditional on any particular coin landing tails it becomes not just more probable, but probabilistically certain that the remaining coins will land tails (and similarly for heads.) This scenario is quite surprising. It involves not just a failure of probabilistic independence but a failure of what I'll call epistemic independence: some epistemically possible combination of outcomes can be ruled out upon learning how an unrelated coin lands – it might be that for all I know the first six coins will land heads, but upon learning how the seventh coin lands (and nothing else) I come to know that the first six coins will not all land heads.

If we combine epistemic independence with the principle that, in ordinary cases, observing the outcome of a particular coin involves acquiring a new piece of knowledge, that the coin landed heads or tails, and does not involve defeating

¹¹This argument is structurally analogous to the puzzle in Dorr, Goodman and Hawthorne [2]. It was this puzzle that first led me to the above observation about evidential probabilities.

any old knowledge, we will be able derive a principle that I shall call CLOSENESS that can be used to show that some of the coins have a probability distinct from $\frac{1}{2}$ of landing tails.

In order to state our principle we need to introduce the concept of one sequence being *closer* to actuality than another. Say that one sequence is at least as close to the actual sequence as another if it agrees with the actual sequence whenever the other one does. More formally, a sequence τ is at least as close as σ to the actual sequence α (written: $\tau \leq \sigma$) if and only if $\tau(n) = \alpha(n)$ whenever $\sigma(n) = \alpha(n)$. It is straightforward to see that \leq is a partial (but not a total) order.¹²

Our principle then states the following:

CLOSENESS: if σ is epistemically possible and τ is closer to the actual sequence than σ , τ is epistemically possible.

In other words K is ‘downwards closed’ with respect to \leq .

CLOSENESS should appeal to epistemologists sympathetic to safety or margin of error constraints on knowledge. A margin of error principle that naturally comes to mind in this context states that a true belief can constitute knowledge only if the belief would have still been true even if a few of the coins had landed differently. If the belief would have been false had a few of the coins landed differently then, although your belief is true, this fact would be too lucky to count as knowledge. ‘A few’ here just means not in excess of a certain number of differences, the number in question being the ‘margin of error’. It’s clear why CLOSENESS is entailed by such a principle – if τ is within the margin (the number of differences from actuality doesn’t exceed the margin of error) and σ is closer to actuality than τ then σ is within the margin too.

A natural route to CLOSENESS, then, is via the margin of error principle outlined above. Unfortunately this would not make for a very compelling argument since the margin of error principle appealed to is at least as controversial as CLOSENESS. Furthermore, even someone who accepts that safety considerations can play a role in determining what one is in a position to know need not accept this particular margin of error principle.

My strategy, therefore, will be to instead prove CLOSENESS from principles that concern how our knowledge evolves when we learn some of the outcomes of the flips. More importantly, principles that do not derive their appeal from safety considerations at all. The two principles I shall be focussing on are:

EPISTEMIC MONOTONICITY: If I observe that a particular coin lands heads (tails) my new knowledge is just the conjunction of my old knowledge with the proposition that the coin landed heads (tails).¹³

¹²Note that this order is isomorphic to a lattice of subsets of 1000 (i.e. sets of three digit numbers.) The reader may find it easier, mathematically speaking, to think of K as a set of subsets of 1000 that is ordered by inclusion.

¹³I think that if we are to have a complete picture of what is going on in these cases we should treat our epistemic state and our evidence as accessibility relations. In this setting monotonicity would state that one’s posterior epistemic state is just the result of intersection

EPISTEMIC INDEPENDENCE: if $X_{i_1} \cap \dots \cap X_{i_n}$ is epistemically possible (where $X_i = T_i$ or H_i) and I discover the outcome of the k th flip where $k \notin \{i_1 \dots i_n\}$ then $X_{i_1} \cap \dots \cap X_{i_n}$ is still epistemically possible.

Epistemic independence is a bit of a mouthful, but it is really very straightforward. It says that if, for all I know, the coins $i_1 \dots i_n$ will land a certain way, finding out how another coin lands will not allow me to come to know that the coins $i_1 \dots i_n$ won't land that way. To all appearances, the outcomes of the coin flips are completely unrelated, and you have no knowledge connecting the outcome of one coin to the outcomes of any others. Suppose I know that five of the coins are on Alpha Centauri and for all I know they'll all land heads. It just seems incredible to think that I could come to know that they won't all land heads by flipping a coin here on earth and observing how it lands. Failures of epistemic independence, however, involve exactly this sort of scenario – a possible outcome of a set of coin flips gets eliminated after learning the result of a distinct and apparently unrelated coin flip.

In the Williamsonian framework evidence is just knowledge. Monotonicity states that if I were to find out tomorrow (by looking say) how one of the coins landed, my knowledge would simply be the conjunction of my knowledge before I looked with the proposition that the coin landed heads or tails, depending on which outcome I observe. I would simply rule out the epistemic possibilities where the coin lands differently from what I've learnt, and nothing that was epistemically impossible before becomes epistemically possible afterwards. Clearly not all cases of knowledge acquisition are like this. Perhaps I can learn something that defeats my earlier knowledge so that things that used to be epistemically impossible are now possible. Perhaps I see that the coin lands heads but I am surrounded by coins that have landed tails and have been painted so as to look like they've landed heads so I fail to come to know what I have (apparently) seen. This is all consistent with the claim that in a great many cases I simply do just come to know a new proposition without undermining any of my prior knowledge. In the coin case, in particular, it is incredibly natural to think that I am in a perfectly good position to know how the coin landed by looking, and that nothing is being undermined, or forgotten, or lost in whatever way by this new knowledge when I do.

The fact that we can acquire knowledge by induction might seem to be in tension with monotonicity principles like the one above. Suppose that I have a coin which I know is either perfectly fair or double headed. It is natural to think that if the coin is in fact double headed, then I can come to know that it is double headed after observing it land heads enough times in a row. Suppose, then, that after learning that the coin landed heads N times in a row (call this proposition A) I am in a position to know that the coin is double headed (call this proposition B .) Thus my knowledge after observing heads N times, K_N , entails B . According to a monotonicity principle structurally analogous to

your prior accessibility relation with the evidence accessibility relation. This has the same effect regarding which worlds are compatible with your knowledge before and after acquiring the evidence.

EPISTEMIC MONOTONICITY, my knowledge after observing N heads in a row is just the conjunction of my knowledge prior to observing any flips, call this K_0 , with the proposition that the first N flips landed heads. Thus $K_0 \wedge A$ entails B , and so by the properties of entailment, K_0 entails $A \supset B$. In other words, prior to investigation I am in a position to know the material conditional: if the coin lands heads N times in a row, it's double headed.¹⁴

For what it's worth, I find the generalized monotonicity principle to provide an extremely attractive model of inductive knowledge. Assuming the coin is in fact double headed and not fair, I'm therefore committed to saying that in favourable circumstances, before even beginning her inquiry, an agent can be in a position to know that if the coin lands heads N times in a row it's a double headed coin. (Of course, whether typical people in fact know, or even believe this, prior to inquiry, is another matter.) Note that this doesn't mean that the agent knows this material conditional *a priori*: the agent wouldn't be in a position to know the conditional in a situation in which the coin is fair but landed heads N times in a row by fluke, for example. Moreover, the alternative picture seems to have the surprising consequence that at the $N - 1$ th flip the following is true: that for all the agent knows the coin is fair and will land heads, but upon observing that it does in fact land heads the agent can knowledgeably conclude that the coin isn't fair after all!¹⁵

At any rate, these considerations are mostly beside the point: the principle I am actually endorsing does not commit us to anything as controversial as this model of inductive knowledge. EPISTEMIC MONOTONICITY applies only to the scenario described in the introduction, and that scenario is importantly different from the above example of knowledge acquired inductively. There is not a single coin, but many, and the chances are known from the beginning: the coins are all known to be fair, and the outcomes are known to be generated from completely independent flips. If we know that the flips are fair and independent, how could learning the outcome of one flip possibly allow us to gain knowledge about the outcome of any other flip.¹⁶ In the inductive case we could only gain knowledge by repeated observation because the possibility that the coin flips weren't fair or independent was open from the beginning!

Finally, even if we ignore the disanalogies with the single coin case, it is totally unclear how this style of response might be applied to a very minor variant of the original puzzle. Imagine that instead of coins, we are flipping

¹⁴I owe this particularly simple presentation of this puzzle to an anonymous reviewer. Note that similar puzzles can be raised for epistemic independence.

¹⁵See also the discussion of the related principle in [2] (there labelled INFERENCE ANTI-DOGMATISM).

¹⁶We can also stipulate that my source for the information about the fairness of the coins is different for each coin. If I had one source who told me they were all fair, and I observed 100 came up heads, I might be reasonable to doubt my source. Whether this means that I'm no longer in a position to know that the coins are fair is more contentious (see Lasonen-Aarnio [1].) But at any rate, it is much harder to push this kind of analogy in the case where one has independent, reliable sources for each your beliefs about the fairness of the coins. This marks an important disanalogy between cases with one coin being flipped many times and cases with many coins being flipped once.

blank metal discs.¹⁷ The puzzle, as I stated it, did not really rely on the fact that it is coins and not metal discs that we are flipping. However, observing a large number of blank discs land according to a particular sequence does not give you evidence that the discs are biased. The fact that this response cannot survive this slight change in the set up should make us skeptical that it can be the source of the problem. Either way, this variant of the puzzle survives and needs an answer that does not rely on this kind of failure of monotonicity. I will continue to talk about coins in the sequel, although it will be useful to refer back to this variant case at various points.

With these two assumptions at hand we can apply them iteratively to tell us what our epistemic state is after learning the outcomes of two or more of the flips – once you’ve learnt the outcome of one flip you will find yourself in a structurally identical set-up (with 999 unknown outcomes instead of 1000.) Monotonicity allows us to infer that after successively learning that coins $i_1 \dots i_n$ land in a certain sequence of heads and tails, X_{i_1}, \dots, X_{i_n} , my new epistemic state should simply be $K \cap X_{i_1} \cap \dots \cap X_{i_n}$. Independence allows me to infer that if $X_{i_1} \cap \dots \cap X_{i_n}$ is epistemically possible (where $X_i = T_i$ or H_i) and I discover the outcomes of coins $j_1 \dots j_k$ where $\{i_1 \dots i_n\} \cap \{j_1 \dots j_k\} = \emptyset$ then $X_{i_1} \cap \dots \cap X_{i_n}$ is still epistemically possible. We can then prove CLOSENESS as follows

Proof. Suppose that $\sigma \in K$ and $\tau \leq \sigma$. Let $S = \{n \mid \tau(n) \neq \alpha(n)\}$ and $\bar{S} = 1000 \setminus S$. Let $X_n = T_n$ if $\tau(n) = T$ and let $X_n = H_n$ otherwise. Note that $\sigma \in \bigcap_{n \in S} X_n \cap K$ so $\bigcap_{n \in S} X_n$ is epistemically possible. It follows by epistemic independence that $\bigcap_{n \in S} X_n$ is possible after learning $\bigcap_{n \in \bar{S}} X_n$, which by monotonicity means it is consistent with $\bigcap_{n \in \bar{S}} X_n \cap K$. Thus $\bigcap_{n \in S} X_n \cap \bigcap_{n \in \bar{S}} X_n \cap K \neq \emptyset$ and since $\bigcap_{n \in S} X_n \cap \bigcap_{n \in \bar{S}} X_n = \bigcap_n X_n = \{\tau\}$ it follows that $\tau \in K$. \square

Let us close this section by proving the desired claim that, given CLOSENESS, for some n , $Q(T_n) \neq \frac{1}{2}$ assuming $1 < |K| < 1000$.

Proof. Let the rank of a sequence, $r(\pi)$, denote the number of places it differs from actuality. Pick a $\sigma \in K$ with maximal rank. If $r(\sigma) = 1000$ then every sequence would be epistemically possible because K is downwards closed. Thus, by hypothesis, $r(\sigma) < 1000$ and it follows that σ matches actuality at some point: for some n , $\sigma(n) = \alpha(n)$. Without loss of generality suppose $\alpha(n) = T$. Note that $\sigma \in T_n$. I claim $|T_n \cap K| > |H_n \cap K|$.

Note that $H_n \cap K$ injects into $T_n \cap K$ by the mapping f that takes a sequence, flips the n th member, and leaves everything else alone. For $\pi \in H_n \cap K$, $f(\pi)$ clearly belongs to T_n . Moreover $f(\pi) \leq \pi$, since π differs from actuality at the n th place and $f(\pi)$ matches actuality at the n th place and otherwise agrees with π . Since K is downwards closed and $\pi \in K$ it follows that $f(\pi) \in K$. So $f(\pi) \in T_n \cap K$.

¹⁷The discs don’t need to have perfectly indistinguishable sides – they might, for example, have distinctive scratches on each side. The crucial point is that there is no way to uniformly classify two discs as having landed ‘the same way up.’ There would be no such way if the scratches on each coin were different.

On the other hand, the sequence $f^{-1}(\sigma)$ (which is, in fact, just $f(\sigma)$) does not belong to K . $r(f^{-1}(\sigma)) = r(\sigma) + 1$ since σ matches actuality at the n th place, and $f^{-1}(\sigma)$ doesn't but is otherwise like σ . Since we chose σ to have maximal rank in K , no element of K has rank larger than $r(\sigma)$ so $f^{-1}(\sigma) \notin K$. Thus $T_n \cap K$ contains at least one more element than $H_n \cap K$ namely, σ . \square

4 Coin Indifference

We have thus far argued that given two plausible claims about how we update in the above cases, the probability that some of the coins will land tails will be distinct from $\frac{1}{2}$. Our motivating question, however, is not yet answered; we know that the relevant probabilities are not always $\frac{1}{2}$, but are they close to $\frac{1}{2}$? And what determines whether the probability is above or below $\frac{1}{2}$?

I will now argue that, for any n , if the n th coin will in fact land heads then your evidential probability that it will land heads is almost 1, and that if the n th coin will in fact land tails your evidential probability that it will land tails is almost 1. In order to show this I will need another assumption, COIN INDIFFERENCE (below), in addition to the assumptions I have made already. Given this assumption we can give an extremely natural class of models for the scenario we have been considering so far.

COIN INDIFFERENCE: The number of epistemically possible outcomes left open after learning how a coin landed does not depend on the coin.

Thus the result of learning how the i th coin lands will leave the same number of sequences epistemically open as learning how the j th coin lands. For example, if I learn how the 27th coin lands and there are 100 sequences which, for all I know, represent the actual outcomes then had I learned instead the outcome of the 121st coin there would have also been 100 epistemically open sequences.

A failure of this principle would leave us with an unexplainable asymmetry between the coins. Some coins would be more informative than others: learning the outcome of some coins would allow us to eliminate more sequences than others. Maybe one could put this down to some asymmetry in the set-up, yet it is clear that one can spell things out to make the coins as symmetrical as you like: the coins could be arranged in a circle in a universe with perfect rotational symmetry and you could know this fact, for example.¹⁸

Even so, one might object, symmetry intuitions are notoriously unreliable in the present setting. According to one compelling symmetry intuition the various different sequences of outcomes are all on an epistemic par (an intuition that becomes clearest when we consider the variant scenario involving blank discs instead of coins.) Yet given our assumptions we know that some of these

¹⁸Kenny Easwaran points out to me that if the universe actually had perfect rotational symmetry we would find out that every coin landed heads after learning that one of the coins landed heads, so we may have to weaken this condition slightly. This issue doesn't arise in the variant puzzle using blank metal discs.

sequences are epistemically possible while others aren't – an apparent violation of symmetry.

This follows from the fact that the actual sequence is epistemically possible (by the factivity of knowledge) and ANTI-SKEPTICISM, which entails that other sequences are not epistemic possibilities. However this reasoning straightforwardly reveals that there is, in fact, a very simple way to break the symmetry between the different possible sequences: if α is the actual sequence, and σ is not, there's an asymmetry right there grounded in the way the coins landed. There need not be anything spooky going on here – the difference between the epistemically possible and impossible sequences can, for all we've said, supervene straightforwardly on the physical arrangement of the coins. This is not so if we were to deny COIN INDIFFERENCE in a perfectly symmetrical universe; there would be an epistemic difference without any physical difference. Surely physicalism is not so easily counterexampld!

In this connection it is useful to distinguish first from third person symmetries. While the possible sequences might seem to be on a par from a first person perspective, they are not on a par from a third person perspective. An onlooker who knew how the coins landed would be able to tell straight away that the actual sequence is epistemically possible for you even if, from your perspective, there is no difference between this sequence and the next. This feature of knowledge should not be surprising – the difference between the epistemic state of a brain in a vat and an ordinary person is not always manifest from a first personal perspective.¹⁹ The kind of symmetry violation that would arise from failures of COIN INDIFFERENCE, on the other hand, would not even be detectable from the third person. Even someone who knew all the physical details about the coins in a symmetrical universe wouldn't have any reason to think any coin more informative than any other.

Once we have accepted COIN INDIFFERENCE, we can iterate the principle so it can be applied to learning the outcomes of several coins in a row. For if I have learnt the outcome of the first coin then I am in a variant of the initial scenario, except that this time there are 999 coin flips with unknown outcomes instead of 1000. COIN INDIFFERENCE, if true at all, is a general truth so we must assume it applies in this variant scenario too. Thus, by successively applying COIN INDIFFERENCE, we know that if the outcomes of all but the coins $\{i_1, \dots, i_n\}$ were learned there would be the same number of epistemically possible sequences as if I were to learn the outcomes of all coins excepts for $\{j_1, \dots, j_n\}$ (provided the i 's are all different, and the j 's are all different.)²⁰

¹⁹Of course, a difference manifests itself in terms of the first personal factive mental states: while an ordinary person is in a position to know that their belief that they have hands constitutes knowledge a brain in a vat isn't. However they are exactly alike regarding their non-factive mental states like belief (and furthermore, their knowledge about what's epistemically possible is identical, even if their knowledge about what isn't differs.) The crucial point is that a third party would be able to have different beliefs about which of the two persons beliefs constituted knowledge (and would know that the 'no hands' proposition was a possibility for one person and not the other.)

²⁰Here is one worry you might have about COIN INDIFFERENCE, particularly the iterated version. Suppose exactly 541 of the coins land heads. Let X denote those coins, and let Y

Once we assume COIN INDIFFERENCE, EPISTEMIC MONOTONICITY and EPISTEMIC INDEPENDENCE we can show that K just consists of those sequences that differ from the actual sequence by at most k places for some fixed k . If K has this form call it a ‘margin for error model’.²¹²²

Margin for error models are strongly reminiscent of a certain safety theoretic way of modelling knowledge. The rough intuition behind safety, as applied to this case, is that one’s belief that p , even if true, cannot constitute knowledge if it could have easily been false. The notion of ‘being easily false’, in this context, is measured by whether the belief would have been false if a few coins (less than k) had landed differently. The intuition is clearer with an example: one’s belief that the coins will not all land heads is usually not only true, but safe in the above sense. However, if by some remarkable coincidence 999 of the coins landed heads, your belief that the coins would all land heads, while true, presumably would not constitute knowledge. If that single coin had landed differently then your belief would have been false: your belief was true by luck and surely could not count as knowledge. Similarly, in cases where all but two of the coins landed heads you surely do not know that not every coin would land heads since it could have so easily been false if those two coins had landed differently. In margin for error models the margin of error, k , represents explicitly how many coins could have landed differently before your belief becomes unsafe.

It is important to remember, however, that we have motivated these models without explicit appeal to safety: the three principles we have appealed to do not appear to derive their plausibility from safety considerations, but rather from considerations about how we update our knowledge. Even if we were to bracket COIN INDIFFERENCE, margin for error models provide a very natural way to model the updating process in these cases in a way that preserves EPISTEMIC MONOTONICITY and EPISTEMIC INDEPENDENCE. I have some ‘accessibility ra-

denote a random set of 541 coins that have roughly as many heads and tails. If you were to have been shown only the outcomes of the coins in X then it might be reasonable to come to believe that all the coins will land heads, whereas it wouldn’t be reasonable to conclude this if you had only observed the outcomes of the coins in Y . Since there would certainly be a different number of doxastic possibilities after making these two sets of observations it’s natural to think the number of epistemic possibilities could also be constrained this way. I think there are three problems with this kind of objection. Firstly, while this thought may have something to it if it were a single coin being flipped repeatedly (for example one might reasonably conclude it was double headed), it is less clear when each of the coins are distinct and isolated in this way. Secondly, whilst the number of sequences compatible with your knowledge does depend on the number of sequences compatible with your beliefs, the number I have been focussing is the number of sequences compatible with what you are in a *position* to know and it is not obvious that this will depend on your beliefs in this way. Finally, this worry cannot even get off the ground in the variant puzzle with blank metal discs since there would be no relevant difference between any pair of sets, X and Y , that have the same cardinality. Thanks to Kenny Easwaran for discussion here.

²¹Models with this structure were first suggested to me by John Hawthorne.

²²Given CLOSENESS it suffices to show that if $\sigma \in K$ and $r(\sigma) = r(\tau)$ then $\tau \in K$. If $\sigma \in K$ then the result of learning the outcomes of all the flips except for those whose outcomes differ from σ will leave us with $2^{r(\sigma)}$ open possibilities by EPISTEMIC MONOTONICITY. If $r(\sigma) = r(\tau)$ then, by iterating COIN INDIFFERENCE, the result of learning the outcomes of all the flips except for those whose outcomes differ from τ will also leave us with $2^{r(\sigma)}$ (i.e. $2^{r(\tau)}$) open possibilities and this can only happen if $\tau \in K$.

dious', k , such that for any n , the possible combinations of epistemically possible outcomes of n flips are those that differ from the actual outcomes by at most k places.²³ This accessibility radius remains constant across time, even as the space of epistemic possibilities shrinks as you learn the outcomes of various coin flips. After finding out the outcomes of the first 500 flips, say, the remaining possibilities are still just those sequences that differ from the remaining 500 actual outcomes by at most k : since you've only ruled out outcomes that are different from actuality (if you've truly come to know or come to bear any factive attitude towards them) we still have all k degrees of freedom left: we can still differ from the actuality in up to k different ways.

Margin for error models, therefore, present us with a particularly simple way of modelling this scenario. Furthermore, as we have argued, to reject this way of modelling the scenario is tantamount to rejecting either EPISTEMIC MONOTONICITY, INDEPENDENCE or COIN INDIFFERENCE.

We can now show that the probability of T_n is greater than $\frac{1}{2}$ whenever the n th coin actually lands tails and $< \frac{1}{2}$ whenever the coin actually lands heads.

Proof. Assume that $n > 1$ and $1 < k < n$. Let $T(n, k)$ represent the number of n length sequences that differ from some fixed sequence by at most k places. Intuitively $T(n, k)$ represents the number of coin sequences that are epistemically possible when you know the coin will be flipped n times, and where k is the margin of error (i.e. k is such that a sequence is epistemically possible iff it differs from the actual sequence by at most k flips. In our case $n = 1000$, but it will be useful to have this extra generality.)

(Aside: mathematically $T(n, k)$ is calculated by the sum $\sum_{r \leq k} \binom{n}{r} = \frac{n!}{r!(n-r)!}$. To quickly perform actual calculations of this number for small n the following trick is useful: draw up the first n rows of Pascal's Triangle and add up the first k numbers on the bottom (n th) row. This fact is not necessary to understand the proof.)

If the first coin (say) actually lands tails, then the number of epistemically possible sequences that begin with tails is just $T(n - 1, k)$, (because once the first element of the sequence is fixed as tails, and that matches the actual world, the remaining $n - 1$ flips can vary from the actual sequence at up to at most k places) so the evidential probability that the first coin lands tails is $\frac{T(n-1, k)}{T(n, k)}$.

Note that for similar reasons the number of epistemically possible sequences beginning with heads is $T(n - 1, k - 1)$ – once the first element of the sequence is fixed, and differs from actuality, then it only has $k - 1$ ways left to differ from actuality.

All that is left is to show that the evidential probability of the first coin landing tails is $> \frac{1}{2}$. This follows from the observation that the number of

²³It is natural to think that 'differing by at most k places' this could serve as an accessibility relation in a model of epistemic logic. This is a reflexive, symmetric, non-transitive relation. These models have the following puzzling feature: out of all of the sequences of outcomes that might for all you know obtain, exactly one of them (the actual sequence of outcomes) is such that you know that it might for all you know obtain. You could consider non-symmetric models where the accessibility radius, k , can vary from world to world.

epistemically possible sequences in which the first coin lands tails (viz. $T(n - 1, k)$) is always larger than the number of epistemically possible sequences in which the first coin lands heads (viz. $T(n - 1, k - 1)$). The former is the cardinality of the $n - 1$ length sequences that differ by at most k places from a fixed sequence, and the former the cardinality of the strict subset of these sequences that differ by at most $k - 1$ places. This completes the proof. \square

Our observation, in the proof, about the cardinality of $T_n \cap K$ in margin for error models provides us with a straightforward way to determine the answer to our opening question: once we know k the probability that a coin will land tails is $\frac{T(999, k)}{T(1000, k)}$ if it will in fact land tails and $\frac{T(999, k-1)}{T(1000, k)}$ if it will in fact land heads. It remains to see what happens when we chose various k .

Which natural numbers, k , might serve as realistic accessibility radii? We can say, straight off the bat, that k is greater than 1. However the coins actually land, we cannot know that the first two coins won't land heads. On the other hand, as we have argued already, k should be less than 1000: we know that the coin won't land heads 1000 times in a row. Could there be circumstances in which we know that the coin wouldn't land heads 10 times in a row? This is, in purely probabilistic terms, equivalent to asking whether one could ever know that their ticket will not win in a fair lottery with 1024 tickets. It seems to me that in ordinary contexts it would be quite reasonable to say that most people do know that they won't win. The case at hand is analogous – it seems reasonable that it's possible for you know that the first ten coins won't all land heads, in which case it follows that k is smaller than or equal to 10. A calculation shows that if $k = 10$ and the first coin actually lands tails the then the probability that it'll land tails is $\frac{T(999, 10)}{T(1000, 10)} = 0.990010172562163$. In other words, if the coin in fact lands tails, we should have an evidential probability of almost 1 that the coin will land tails! A completely analogous argument applies if the coin in fact lands heads.

Of course, you could quibble about my choice of k . Quibbling makes little difference to my point. Here are the resulting probabilities, to two decimal places, if we substitute k for larger numbers²⁴:

k	$Q(T_n)$
10	.99
20	.98
30	.97
40	.96
50	.95
60	.94

To get an idea for these numbers, note that any model in which $k \geq 60$ is one in which for all you know the first 60 coins will all land heads. Not knowing whether

²⁴These numbers were calculated using <http://www.wolframalpha.com/>. To calculate $Q(T_n)$ when $k = 10$, for example, just input: `sum[binomial(999, r), {r, 0, 10}]/sum[binomial(1000, r), {r, 0, 10}]` (note: it seems to crash when you chose $k \geq 70$.)

the coin will land heads 60 times in a row is equivalent, at least statistically speaking, to not knowing whether your ticket will lose in a lottery with one quintillion (i.e. one billion billion) tickets.²⁵

5 Consequences for the Lottery Paradox

Let us, for a moment, set aside the problem for evidential probabilities. In this paper I have not just argued for a negative thesis, but for a positive view about the anatomy of our knowledge in lottery paradoxes with this particular structure. It would be worth our while to compare the resultant margin of error model of knowledge to other accounts of the lottery paradox. If these alternative accounts do not fair as well as the safety model in lotteries with this particular structure, then this might point favourably to a safety based account of lottery paradoxes in general.²⁶

Central to all these puzzles is the surprising result that if one accepts the anti-skeptical idea that we are in a position to know that some sequences will not occur, then it follows that not all of the possible sequences are on an epistemic par. There will be some sequences that our knowledge rules out, and other sequences which it doesn't: factivity guarantees that at least one sequence – the actual sequence – is epistemically possible, and ANTI-SKEPTICISM guarantees that at least one sequence isn't. However, from a first personal perspective no sequence seems different from the next. How, one might wonder, could one be in a position to rule out some sequences but not others?

Fortunately philosophers have become adept at resisting this style of questioning: from a first personal perspective there is a clear symmetry between myself and a brain in a vat, but unless we are to surrender ourselves to skepticism we must concede that I'm in a position to rule out possibilities that the brain in a vat isn't (for example, I'm able to rule out the possibility that I don't have hands.) However, in the brain in a vat case it is clear which external facts ground the epistemic asymmetry; I have hands, for example, but the brain in a vat doesn't. Given that there are also epistemic asymmetries between the different sequences it is consequently very natural to wonder which facts – also, presumably, external to the agent – the asymmetries in our present case depend on. I have defended the view that your knowledge in these cases depends on how the coins in fact land, and the value of a certain margin for error; it is thus

²⁵This observation seems to be pretty stable provided k does not depend on the number of coins in the set up: whatever k is, you can also modify the number of coins to be flipped (currently 1000) appropriately to achieve a similar result. One could in principle avoid this result by allowing the structure of knowledge to depend on things like the number of coins flips to occur throughout all history, or on how many of those flips with unknown outcomes you know to be fair, but these would be drastic measures.

²⁶Articulating a measure of closeness to actuality relevant for an arbitrary lottery, without knowing the mechanism by which the winner is generated, is a tricky business. One would also expect there to be a certain degree of context sensitivity regarding which possible outcomes to regard as closer to actuality than others. However, even if one cannot give a completely explicit account of closeness, safety based theories will still impose structural constraints on knowledge that will distinguish them from other approaches to the lottery paradox.

clear how the symmetry is broken on this picture. I'll now compare this idea to three alternative accounts:²⁷

1. No sequences are ruled out: there are no sequences we know won't happen. Widespread skepticism is quarantined by restricting single premise closure.
2. Abnormal sequences are ruled out: surprising sequences like the all heads or all tails sequence are ruled out, unsurprising mixed sequences are not ruled out. The asymmetries are grounded in facts about which sequences are surprising.
3. Non-actual sequences are ruled out: each sequence, apart from the actual sequence, is known not to occur. The only asymmetry is between the actual and the non-actual sequences, and this is grounded in the way the coins land.

According to the first kind of theory all the sequences are apparently on an epistemic par with the actual sequence: it is not known whether they will obtain or not. On this picture it seems as though there are no asymmetries to be explained away. However, if we are to resist skepticism, there must be some sequences that are 'ruled out' by my knowledge in an extended sense: I know things (e.g. that I won't have enough money to go to the Bahamas next year) that entail that I won't win the lottery. So at least in some sense there is still an asymmetry to be explained. At any rate, we can all agree that the epistemic difference between the sequences isn't as direct.

I think there is much to be said for this kind of view. Many people, myself included, have the intuition that typical people do not know that their (or perhaps any) lottery ticket won't win. Indeed, I expect if you look at the average person taking part in a lottery that intuition is true: they probably wouldn't buy a lottery ticket in the first place if they knew that it wasn't going to win. However, there are plenty of other things they do know that entail that their ticket won't win. As this theorist predicts, knowledge doesn't appear to be closed under logical consequence.

One might think I am being overly concessive to this kind of theorist, but I suspect that all I have said above is common sense. The reason people buy lottery tickets is because they don't really believe with enough conviction that they'll lose. Moreover, belief isn't an attitude that is closed under logical consequence – there are plenty of propositions that follow from our beliefs that we do not believe because we have never even considered them. Any attempt to formally model knowledge must be sensitive to the fact that one can fail to know something by simply not believing it, and which propositions a person believes in a given lottery case depends on all kinds of contingent facts about

²⁷Note that there is also a recent tradition of taking practical factors (such as how much is at stake if you were to act as if p) to impact what you know. Theories in this vein will not be always compatible with some of the principles I appealed to in this paper (see COIN INDIFFERENCE and EPISTEMIC INDEPENDENCE.) I take this to be a problem for these theories rather than a problem for the principles. However, for the sake of space, I shall simply bracket these views from consideration.

the person. A practising skeptic might refuse to believe, and thus fail to know, anything about the outcomes of the coins whilst a person less skeptically inclined might know that not all of the coins will land heads, although even the non-skeptic won't know things about complicated sequences of outcomes she's never even entertained.

If we are to say anything theoretical here we must distinguish between what a person knows, and what she is in a *position* to know. In the preceding example the skeptic and the non-skeptic are on a par regarding what they are in a position to know, the difference is whether they in fact believe those things or not.²⁸ The more important question, for our purposes, is: if I know that I won't have enough money to go to the Bahamas next year, am I in a position to know that I won't win the lottery? It is much harder to maintain that a person who has performed the relevant deduction is not in a position to know this. So while our intuitions about our lack of knowledge in the lottery paradox may remain intact, it is entirely possible that we are in position to know a reasonable amount about which tickets won't win.

Assuming, then, that we accept ANTI-SKEPTICISM it is natural to ask which sequences are ruled out? While ANTI-SKEPTICISM doesn't commit us to any particular sequence being ruled out, the all heads sequence seems like a promising candidate. But what is it about the all heads sequence that means it can be ruled out over the other non-actual sequences? A natural thought is that it is surprising, or unusual, unlike the random looking sequences you'd expect to see most of the time.

This is how the second type of view breaks the symmetry: some sequences are more surprising or 'unusual' than others, and it is this feature that the all heads sequence has that distinguishes it from the sequences we cannot rule out. It is important to note that 'surprising' in this context doesn't mean unlikely, or atypical: each possible sequence is just as likely as any other to occur, and no particular sequence will occur more often than any other on average, so they are all equally typical.

As we have seen already, if surprise or unusualness can be a source of epistemic asymmetry then there are natural ways to reject some of the principles I have discussed earlier. One might motivate unusualness as a source of epistemic asymmetry by noting that, upon observing a coin land heads 100 times in a row it seems rational to conclude that it's not a fair coin, or is in some other way unable to deviate from this sequence, but upon observing that it lands in some random unsurprising sequence featuring roughly as many heads as tails it would not be rational to conclude that it wasn't fair.²⁹ The idea that it would be rational to conclude that a set of coins aren't fair after observing 100 of them

²⁸My rough gloss of this notion is: one is in a position to know that p if and only if there is some way of coming to believe that p that would constitute knowledge.

²⁹Thanks to Jeremy Goodman for this point. This thought is probably behind many supporters of 'intelligent design': that one can knowledgeably conclude that some sequence of outcomes is not the result of chance if the results are sufficiently ordered or surprising, even if that sequence is just as probable, conditional on being chancy, as a particular unsurprising random sequence of outcomes.

land heads, but not rational after observing them land according to a particular seemingly random sequence, if plausible at all, must rest on facts peculiar to coins (for example, that a salient way for the coins to be unfair is for them to be double headed.)

It is important to emphasize, then, that analogous puzzles can be formulated in a way that does not depend on whether unusualness is a source of epistemic asymmetry. What makes a sequence of coin flips surprising or unsurprising depends crucially on the fact that coins have been marked in a certain uniform way to distinguish between the two sides. If instead of coins we used, say, blank metal discs then no sequence would be more surprising or unusual than any other; in this variant scenario unusualness clearly cannot be contributing to determining which sequences are epistemically possible.³⁰ Note, however, that in the blank metal disc variant it seems like the only real source of asymmetry between the sequences is an asymmetry between how they actually land, lending support to my hypothesis that this is the only source of asymmetry in the coin case.

According to the third alternative you are in a position to know of each of the non-actual sequences that it won't obtain. Thus there is as much symmetry as there can be given that the actual sequence isn't ruled out and that some sequences are ruled out. The asymmetry between the actual and non-actual sequences is easily grounded in the way that the coins actually land. This view also has a lot of appeal, but it is subject to a couple of problems.

The first of which is the well-known point that it has absurd consequences when combined with multi-premise closure. If one could close one's knowledge under conjunction introduction you would be in a position to find out how the coins landed without looking at them; but surely that's absurd, nobody can know this *a priori*.

The second problem arises specifically for lotteries that are generated by sequences of independent events, such as in our present example. The problem is that this model conflicts with a highly desirable margin for error principle:

MARGIN FOR ERROR: If τ describes the actual outcome of the coin flips, then for all you're in a position to know, the coins landed according to τ' , where τ and τ' are two sequences of outcomes that differ from one another over the outcome of one flip.

The intuition behind the margin for error principle is clear, and follows from a much more general intuition: Suppose that tomorrow a coin will be flipped and will land tails. Now, if p is such that it would have been false if the coin had landed heads, then it surely follows that a belief that p , even if true, could not constitute knowledge. I would still have come to believe that p via the same

³⁰It doesn't strictly speaking matter that they be blank discs. One might worry that it is impossible to assign distinct epistemic possibilities to the different ways the coin can land unless one can distinguish the sides. For the record I disagree, but it does not matter: one can assume that the sides are distinguishable, but not in a uniform way like with coins (for example, perhaps the first disc has a scratch on one side, the second disc has different coloured sides, and so on.)

method had that coin landed heads, so it follows that I could easily have falsely believed that p .

The above margin for error principle can be seen to be a particular instance of this type of reasoning. Suppose that by coincidence the first 999 of the coins land heads, and the final coin lands tails. In this scenario it is absurd to suppose that I knew that the coins wouldn't all land heads. If that last coin had landed differently then my belief would have been false – I was lucky that it was even true. Yet if we are to apply the third model of the lottery paradox listed above to this particular paradox it would follow that I *was* in a position to know that the all heads sequences wouldn't happen.³¹

Contrast these accounts with my preferred margin of error model. Like the third model, the asymmetry between the epistemically possible and impossible sequences is grounded by the way the coins actually land (once you've fixed the value of the margin of error.) However, this model is compatible both with multi-premise closure and the margin of error principle. Unlike the second model, it can be applied straightforwardly to the variant puzzle involving blank metal discs. Furthermore, we have argued that denying this model is tantamount to denying either EPISTEMIC MONOTONICITY, EPISTEMIC INDEPENDENCE, COIN INDIFFERENCE or the closure assumptions that are implicit in this kind of model.

Despite these credentials, there are two features of these models in particular that require some comment.

The first feature is that the coin flips are not pairwise probabilistically independent of one another, which is to say, the outcome of one coin flip can evidentially support hypotheses concerning the outcome of another coin flip. This is of course assuming that probabilities are Williamsonian evidential probabilities; other notions of probability can be introduced for which this result doesn't apply. It is not obvious that we can really avoid these kinds of results: we have proved already, without making any assumptions other than ANTI-SKEPTICISM and THE PRINCIPAL PRINCIPLE, that the coins are not *mutually* independent of one another (which is to say the outcomes of a *collection* of coin flips can be made more or less likely by the outcome of a distinct coin flip.) Indeed, it is very hard to see how to construct a model that guarantees that the coins are pairwise independent of one another given the fact that they are not jointly independent.

The other feature of these models requires a bit more discussion. To simplify things, let us suppose that that $k = 10$. According to the margin for error model, in worlds where actual outcomes of the coin flips are given by σ , K will just be the set of sequences in which at least 990 of the outcomes match σ . Let's suppose, for simplicity, that all of the coins land heads: then K is simply the set of sequences in which at least 990 of the coins land heads. If we are not careful about how to interpret this, this observation would seem to suggest the absurd claim that, in the cases when the coins all land heads, I know that at

³¹This objection also applies to a naïve version of the second model. A more sophisticated version of the second model might set things up so that the more unusual things are, the less you know (see the ideas described Goodman [3] section III, which is based Williamson's analysis of Gettier cases from [9].)

least 990 of the coins landed heads!

It is important, when considering these puzzles, to recall our earlier comments to the effect that K must, if we are to do any theorising, represent not what an agent *knows* but rather what she is in a *position* to know. I take it that when the coins in fact land heads 1000 times the average person will not have the belief that the coin will land heads at least 990 times – a very surprising thing to believe – and therefore will not in fact know this even if she is in a position to.

Nonetheless, there is still a nagging worry. Perhaps it is true that, in these circumstances, you are in a position to know that 990 of the coins will land heads. But the fact that this is something that you are in a position to know depends on how the coins actually landed, and in particular, on the fact that the coins all landed heads (if they had landed any other way you would not be in a position to know that 990 of the coins landed heads.) Since I don't have access to the facts about how the coins landed, I don't have the information I'd need to determine what I'm in a position to know. How would a person even go about working out which sequences to *believe* won't be actual? Even if the sequences aren't all on an epistemic par, the sequences are all on a par from the first person perspective of the agent.

Of course, one might insist that the standards by which we judge peoples beliefs to be doing well epistemically depend on these external conditions as well: an agent who believes everything that she is in a position to know will be doing well epistemically, and the conditions for knowledge are externalist ones. But still, the worry is that one cannot actively achieve this state except by fluke, and it would not do to theorise about an epistemic state that one could only find oneself in by fluke.

It would be desirable to have a strategy, or some advice, that one can follow and know that one is following – call this an 'operational' strategy – that will result in one knowing everything one is in a position to know whenever the strategy is followed. Luckily, I think, there is at least one clearly operational strategy that one can follow: for each sequence τ , simply believe that τ will not happen. Things being favourable your beliefs in the propositions that you are in a position to know will constitute knowledge; you will also have some beliefs which do not constitute knowledge and you will have exactly one false belief. However, these bad cases are vastly outnumbered by our beliefs that constitute knowledge (if $k = 10$ then the fraction of these beliefs that aren't knowledge will be tiny: $\frac{T(1000,10)}{2^{1000}}$).

Note that the strategy does not recommend believing arbitrary conjunctions of these propositions. Indeed, the conjunction of these beliefs is inconsistent so one cannot combine this belief forming strategy with the 'conjunction introduction' strategy for forming beliefs. I must stress here that I am not suggesting that this strategy, and not conjunction introduction, is the strategy one ought to follow. A natural thought would be that what one *ought* to do is believe exactly what one is in a position to know – both the aforementioned strategy and conjunction introduction are followable rules that can lead to this goal, but

I doubt there is any (followable) rule that will tell you when to use which.³²

6 Concluding Remarks

Might the results in this paper be generalised beyond the fairly idealized setting assumed throughout this thought experiment? For example, suppose you were simply to take a coin out of your pocket and flip it. What might we conclude about the evidential probability of it landing heads before it's flipped, supposing it in fact does land heads? If the results of this paper are at all suggestive one might expect the answer to be a very high number, and certainly not a half. While this conclusion relied on the presence of the other 999 coins in the thought experiment, it seems that formally all one needs to be able to do is divide a region of logical space in which we live into a partition of equal chances regarding the outcomes of a large sequence of seemingly independent events.³³

What should we conclude from all of this? Williamson informally introduces evidential probabilities in terms of the kinds of answers we'd give to questions about how probable this or that proposition is on the body of evidence we currently possess. Yet it seems fairly natural to think that the heads outcome is not extremely probable given our present evidence, *pace* Williamson's definition of evidential probability in terms of knowledge. Perhaps this is just to beg the question against Williamson's hypothesis that evidence is knowledge.

A better way to approach this question is to ask what roles evidential probabilities can play given they have these features. I think one very straightforward conclusion would be that evidential probabilities cannot correspond to rational betting odds: it seems clearly irrational to bet large amounts of money on a coin you know to be fair landing heads (even if you do end up winning in the cases where the evidential probability is high.) Intuitively it's permissible to accept at most even odds when offered a bet on the outcome of a coin you know to be fair.

Evidential probabilities do not seem like a particularly good measure of the justification that a proposition has either. If all I know is that the coin is fair, I have no more justification for the proposition that it will land heads than for the proposition that it will land tails; my level of justification has nothing to do with the way the coin actually lands.

The problem, I think, is that in lottery cases one typically knows a lot more than what one strictly speaking has evidence for. You may know that not every

³²A natural way to model belief on this operational understanding would be to use sets of propositions that are represented by some probability/threshold pair: that for some Pr and α you believe that p iff $Pr(p) > \alpha$ (there is a natural relation between k and α .) Your beliefs would be closed under single-premise closure, but not multi-premise closure. In this model they would also be closed under a recombination principle, which states that whenever you believe $A_1 \dots A_n$, and when $P_1 \dots P_n$ is a partition and $A^* = \{X_1 \cap \dots \cap X_n \mid X_i = A_i \text{ or } \neg A_i\}$, then there is a function $g : A^* \rightarrow \mathcal{P}(P)$ such that (i) $|g(X)| \leq |\{A_i \mid X \subseteq A_i\}|$ and (ii) you believe the proposition $\bigcup_{X \in A^*} (X \cap (\bigcup g(X)))$.

³³Certainly there have been at least 1000 coin flips throughout the course of history whose outcomes are unknown to us, thus it seems natural to think that such a partition ought already to be available to us.

coin landed heads even if you haven't seen, remembered or born any other kind of evidential attitude towards the proposition that not every coin landed heads. A very natural constraint, that Williamson's theory violates, is the requirement that your total evidence at time t be admissible at t . In fact, in conjunction with THE PRINCIPAL PRINCIPLE, this principle entails that your evidential probability matches the chances in the way expected.

It would be reasonable to object that the notion of 'admissible evidence' is too unclear to be doing any important theoretical work – while Lewis's two examples of admissible evidence provide a good rule of thumb, the idea that one can sharply distinguish between propositions that are about matters of particular fact no later than the present from similar propositions about local matters in the future rests on a controversial Humean metaphysics. Luckily these issues can simply be sidestepped by replacing the notion of being admissible at t with the notion of having a chance of 1 at t , conditional on any proposition.³⁴ (Those who do not have any reductionist ambitions regarding the notion of chance can simply think of admissibility this way, but this goes far beyond my purposes.) This suggests the following constraint on evidence:

THE CHANCE-EVIDENCE HYPOTHESIS: If E is part of your evidence at time t , E has a chance of 1 at t , conditional on any proposition.

Theories that violate this hypothesis, like Williamson's, are potentially susceptible to puzzles of the kind I have raised here so it is only natural to consider what happens if we impose this restriction.

In this regard it is worth noting that it is possible to formulate a theory of evidence satisfying this hypothesis that still accords with some of the basic insights of Williamson's theory. An important aspect of that theory is that it characterises evidence by its interaction with a certain class of factive propositional attitudes. For Williamson one's evidence is coextensive with the disjunction of all these factive mental state operators³⁵, including knowledge, and it is this feature that causes trouble. A more moderate theory might characterise evidence by its equivalence with the disjunction of some more limited set of propositional attitudes that correspond to more recognizable ways of acquiring evidence. At minimum I would include in this list attitudes expressed by the expressions ' S could hear that p ', ' S remembered that p ', ' S saw that p ', ' S could feel that p ', (see [8] for a more comprehensive list of these operators, with accompanying discussion.) However it would not include ' S knows that p ' – in the case described here, even though one knows the coins will not all land heads, one strictly speaking does not have evidence that they will not all land heads for this is something one has neither seen, heard, remembered nor is it something one has stood in any of these more concrete evidential relations to.

³⁴The addition of the clause 'conditional on any proposition' is redundant if we simply take conditional probabilities to be defined from unconditional probabilities (although, for technical reasons, it would be inadvisable to adopt such a definition.)

³⁵By the disjunction of two attitude verbs, V and V' , I mean the complex attitude that obtains when one V 's that p or V' 's that p .

References

- [1] Maria Lasonen Aarnio. Unreasonable knowledge. *Philosophical Perspectives*, 24(1):1–21, 2010.
- [2] Cian Dorr, Jeremy Goodman, and John Hawthorne. Knowing against the odds. *MS*, 2013.
- [3] Jeremy Goodman. Inexact knowledge without improbable knowing. *Inquiry*, 56(1):30–53, 2013.
- [4] John Hawthorne. *Knowledge and lotteries*. Clarendon Oxford, 2004.
- [5] John Hawthorne and Maria Lasonen-Aarnio. Knowledge and objective chance. *Williamson on knowledge*, pages 92–108, 2009.
- [6] David Lewis. Philosophical papers, volume ii. *New York: Oxford University*, 1986.
- [7] Jonathan Vogel. Are there counterexamples to the closure principle? *Philosophical Studies*, 48:13–27, 1990.
- [8] Timothy Williamson. Knowledge and its limits, 2000.
- [9] Timothy Williamson. Gettier cases in epistemic logic. *Inquiry*, 56(1):1–14, 2013.