AN ETERNAL CON

RANDALL G. MCCUTCHEON

ABSTRACT. Cian Dorr considered the case of a fair coin that is tossed every day throughout an infinite past and an infinite future (an "Eternal Coin"). Against intuition, he argued that, conditional on the Coin having landed *heads* throughout the past, one should believe, with full probability, that it will also land *heads* today. In this paper, we critique Dorr's arguments, as well as part of a reply by Myrvold.

1. The Eternal Coin, Revisited

Cian Dorr (2010) invites us to consider an "Eternal Coin".[1] This is a coin that has been tossed every day throughout an infinite past, will be tossed today and will continue to be tossed, once per day, throughout an infinite future. The Coin is fair, the tosses are independent (everything to come is conditional on these facts) and we know that we will never have any information about the outcomes of the tosses. Consider the propositions

H: The Coin lands *heads* today.
P: The Coin landed *heads* on every past day.
F: The Coin will land *heads* on every future day.

Dorr argues that the rational conditional credences to hold are $Cr(H|P) = Cr(H|F) = 1$. (Call this position ONE.) Our purpose is to show that Dorr's arguments do nothing to defeat the presumption that $Cr(H|P) = Cr(H|F) = \frac{1}{2}$. (Call this position HALF.) A further question is whether HALF is worth adopting, or if rather these conditional credences are better left undefined. (Undefined is cleaner, though for the related problem of conditionalization on the most natural σ-algebra, we will say $\frac{1}{2}$.)

First we argue against the conclusion of ONE directly. Suppose there are two independent Eternal Coins (Coins A and B), and a third coin (Coin C). Each day, Coin C is placed in the *heads* position if and only if the outcomes of that day's Coin A and Coin B tosses coincide. Coin C's outcomes constitute a $\frac{1}{2} - \frac{1}{2}$ *IID* (as with the original Coin, the random variables governing the placements are *independent* and *identically distributed*) process. Accordingly, an advocate for ONE is committed to

[1]This ill-fated paper is an essentially unedited amalgamation of an initial submission and revision written for *Erkenntnis* in 2017. "Ill-fated" not because it was rejected, but because I had naively assumed (see footnote 8 below) that Dorr was implicitly conditioning on the most natural σ-algebra, and I had believed (as I still do) that the answer to *that* puzzle should be $\frac{1}{2}$ (in particular defined, contra Myrvold 2015). Appraised of Dorr's true intention by referees, I soon realized that the argument at the end of the text below (cf. Builes 2019, where the same argument is presented with different emphasis) commits the $\frac{1}{2}$ solution to the problem Dorr did in fact pose to mere finite additivity (and so isn't acceptable). A thematically consistent revision was therefore impossible.

holding that, conditional on Coin C having been placed *heads* throughout the past, it will almost surely be placed *heads* today.

This result is not comforting. For as Dorr acknowledges, there is good enough reason to think that one should have full credence in the proposition that the actual past σ (a function from the negative integers to $\{0,1\}$; $\sigma(-n)=1$ if and only if the coin in question landed *heads* n days ago) of Coin A is such that, conditional on it, one would have credence $\frac{1}{2}$ in Coin A's landing *heads* today. So if the ONE advocate were to first learn that the pasts σ_A and σ_B were the same (thus coming to believe, with full probability, that Coins A and B coincide today), then were to next learn the common value of the pasts $\sigma_A=\sigma_B$, she would almost surely (by Reflection, say) have come to believe that two independent, $\frac{1}{2}-\frac{1}{2}$ Bernoulli random variables (those governing today's A and B tosses) agree almost surely. Given the transparency of the situation (self-locating effects, if any, ought to have manifested in the individual Bernoulli variables), that is absurd.

We turn now to Dorr's arguments. Consider first an *argument from self-locating reference*:

> Take any jointly inconsistent centred propositions A and B that agree about what the world is like as a whole and who you are, and disagree only about when it is. Suppose that if either A and B is ever true of you, each must be true of you for the same length of time - e.g. one day. And suppose that your evidence does nothing to discriminate between A and B, in the sense that *a priori* you should regard it as no more or less likely given A than given B. Then you should regard A and B as equally likely, in the strong sense: conditional on $A \vee B$, your credences in A and in B should be equal.

Dorr now lets K_0 be the centered proposition *today is the last Tails day preceding an infinite final run of Heads*, and lets K_i be the centered proposition *today is the ith day of an infinite final run of Heads*, $i=1,2,\ldots$. By the above indifference principle,

$$Cr(K_0|K_0 \vee K_1 \vee \cdots \vee K_n) = \frac{1}{n+1}$$

for each n, from which it follows that $Cr(K_0|K_0 \vee K_1 \vee \cdots) = 0$. And from here, of course, it is not far to the conclusion that $Cr(H|F)=1$.

We claim that the argument from self-locating reference overgeneralizes. Let F_π be the proposition that the future of the Coin is equal to some tail of the binary expansion of π (with 1 for *heads* and 0 for *tails*), and let F'_π be the proposition that the present and the future of the Coin taken together are equal to some tail of the binary expansion of π. Let L_0 be the centered proposition *today is the last day preceding a maximal infinite final run equal to some tail of the binary expansion of* π, and let L_i be the centered proposition *today is the ith day of a maximal infinite final run equal to some tail of the binary expansion of* π, $i=1,2,\ldots$. Again by the indifference principle,

$$Cr(L_0|L_0 \vee L_1 \vee \cdots \vee L_n) = \frac{1}{n+1}$$

for each n, from which it follows that $Cr(L_0|L_0 \vee L_1 \vee \cdots) = 0$, implying that $Cr(F'_\pi|F_\pi) = 1$. In other words, the Coin is to be expected to eerily continue (into the past) any serendipitous mirroring of π. This appears to be in tension with Dorr's later claim that for a "typical past", with "no interesting periodic behaviour", there is "nothing to defeat the presumption" that rational credence in *heads* ought to be $\frac{1}{2}$, conditional on that past.

Moreover, were one to first learn F_π (and so come to have credence 1 in F'_π), and then learn the future trajectory of the Coin, one would, assuming one were well enough acquainted with the binary expansion of π, almost surely (Reflection, again) have come to possess an extreme credence in the outcome of today's toss. This is especially strange in that one could at this point reflect on the fact that there is nothing special about π here; one might well have started the thought experiment with some number having a binary expansion last differing from that of π at the place corresponding to today's toss.

Dorr considers a counterproposal, according to which one should affirm

$$Cr(K_{n+k}|K_{n+k} \vee K_n) = \frac{1}{2^k + 1},$$

but neglects drawing out its more compelling features. A strong case can be made by considering how K_i constrains tosses in the vicinity of today's. K_1, for example, constrains yesterday's toss (*tails*), today's toss (*heads*), and every future toss (*heads*). K_2 also constrains yesterday's toss (*heads*), today's toss (*heads*), and every future toss (*heads*), but it also constrains the day before yesterday's toss (*tails*). So there is plausible reason to deem K_2 "one-half as likely as K_1, conditional on their disjunction".

Next consider an *argument from evolving credences*. Let Cr denote the credence function you have today, and let Cr_+ be the credence function you will have tomorrow. Dorr accepts the following as premises:

a. $Cr_+(HF|P \vee HF) = Cr(F|HP \vee F)$
b. $Cr_+(P|P \vee HF) = Cr(HP|HP \vee F)$
c. $Cr_+(HF|P \vee HF) = Cr(HF|P \vee HF)$
d. $Cr_+(P|P \vee HF) = Cr(P|P \vee HF)$
e. $Cr(P|P \vee F) > 0$
f. $Cr(F|P \vee F) > 0$

Dorr refers to premises a. and b. as UPDATING, to premises c. and d. as CONSTANCY and to premises e. and f. as POSITIVITY. From them (as well as some innocuous assumptions about the algebra of conditional probabilities), he derives $Cr(H|P) = Cr(H|F) = 1$.

In a "preliminary defense" of the CONSTANCY premises, Dorr asserts that since one's total evidence "does not change in any relevant way between today and tomorrow", one's conditional credences in centered propositions about the Coin should not change either. To the charge that mere passage of time constitutes relevant evidence, he offers the following:

> ...in an ordinary case where the passage of time diminishes your credence in the centred proposition *it is raining*, your evidence today may include the centred proposition *it is raining right now*...while your evidence tomorrow instead includes the centred proposition *it was raining yesterday*.... Nothing like this happens with the Coin.

This comparison is misleading. *It is raining* is plainly analogous to *the Coin lands heads today*, not to *the Coin has always landed heads*. The latter is analogous to *it has always been raining*, and there isn't any reason to believe that the passage of time should be irrelevant to one's conditional credences about that centered proposition in a case where rain behavior is stipulated to arise from an IID process.

In a "further defense" of the CONSTANCY premises, Dorr paints the following picture of one aspect of HALF reasoning:

> As time goes on, you regard HF as more and more likely relative to P. (...) If the ratio of your credences in HF and P, conditional on their disjunction, is x today, tomorrow it will be $4x$. (...) You spend almost all of your life regarding one of P and HF as vastly more likely than the other.

It is intended that this narrative be damning, but we think it's potentially more the opposite, not least because if the Coin were tossed a very large finite number of times and one's (much shorter, say) lifespan were to be situated uniformly at random inside that interval, that (or something very close to that) is just how one's conditional credences would in fact evolve.

From the initial defense of the POSITIVITY premises, meanwhile: "For POSITIVITY to fail would be for you to regard one of F and P as infinitely more likely than the other. Such a bias seems odd, given the symmetry of the setup." But the more promising option is to leave the conditional credences $Cr(P|P \vee F)$ and $Cr(F|P \vee F)$ undefined. Indeed, there seems to us to be good reason for doing so.[2] Dorr writes:

> Compare the present case to one where infinitely many Coins are tossed simultaneously, in an infinite row of houses running East-West. In that case, it would be crazy to regard *Heads in every house to the East of me* as infinitely more or less likely than *Heads in every house to the West of me.* Why should it be any different when the distribution is temporal rather than spatial?

[2]Noting that HALF defenders must allow some instances of conditionalization on null events, Dorr writes "...perhaps there is something special...which prevents $Cr(P|P \vee F)$...from being well-defined." One might point to chances, which feature so crucially in the argument for HALF. There doesn't, to us, seem to be a fact of the matter about the relative chances of P and F, conditional on their disjunction, whereas the chance of H conditional on P is $\frac{1}{2}$ by stipulation. (The tosses are independent and the Coin is fair.) But even if there are non-trivial chances, then surely they evolve in the "counterintuitive" way envisioned by Dorr in his argument for CONSTANCY, so that there would be a unique first day at which $|Ch(P|P \vee F) - \frac{1}{2}|$ achieved its minimum value. Saddled then with merely finitely additive credences over the possible values of one's displacement from that unique day, we wouldn't know how to compute the expectation of these chances.

Suppose one regards “Heads East” and “Heads West” as equally likely conditional on their disjunction. One concern is what would happen if one were to now (say) burn down every second house to the East, put the remaining houses on trucks to fill in the gaps, and toss again. On the one hand the picture prior to the second toss looks just like the picture prior to the first toss; there are infinitely many houses to the East and infinitely many houses to the West. So perhaps one ought to regard “Heads East, 2nd toss” and “Heads West, 2nd toss” as equally likely, conditional on their disjunction. On the other hand, the Coins in the houses to the East that are being tossed the second time are precisely those that were in odd numbered (say) houses to the East the first time. So perhaps one ought to regard “Heads East, 2nd toss” as infinitely more likely than “Heads West, 2nd toss”, conditional on their disjunction. (As one is already committed to regarding “Heads East, odd house numbers” as infinitely more likely than “Heads West”, conditional on their disjunction.)

The closest thing to a “smoking gun” passage in Dorr (2010) is this, in which are contemplated worlds that contain, in addition to an Eternal Coin, a gong that rings just once.

> There are at least two different ways to set up a conditional chance function over such worlds, each of which has a claim to represent the Coin as fair, the gong as equally likely to ring on any given day, and the two as independent. Roughly speaking, we could either start with a conditional probability function over propositions about how the Coin lands each day and then fine-grain its domain to include propositions about the gong, or else we could start with a conditional probability function over propositions about when the gong rings and then fine-grain its domain to include propositions about the Coin.

To start, we don’t think that Dorr is sufficiently wary of the problems caused by the violations of countable additivity brought on by one’s own displacement from the ringing of such a gong. To bring them into relief, suppose there are two such gongs and that you are offered your choice of them, winning 2 dollars if the ringing of the one you choose has greater (in magnitude) displacement than the other from your current position. Once you choose, it’s understood that you will learn the displacement of the ringing of the gong you chose from your current position and be offered the opportunity to switch...at a cost of one dollar. Since upon learning the displacement of the gong you chose from your current position you will come to have credence 1 that the other’s ring will have a greater magnitude displacement from your current position, you will definitely pay the dollar to switch. But you know all of this in advance, so you will just keep your dollar and choose the other gong in the first place. Until, that is, you realize that the same argument blocks the choosing of that gong as well.

We think there is a way to run the chance function debate in such a way that at least one of the two options (the stronger, we think) avoids any run-in with mere finite additivity. So far as we can tell, the gong serves two purposes. The first is to provide (in the case of the “gong first” chance function) a center of reference around which to build a Coin measure. E.g. *heads the day the gong rings* gets measure one-half, *heads the day the gong rings and tails the day after* gets measure one-quarter, etc. That purpose can easily be served (indeed usually is served) by some designated position defined in relation to one's current location, so that *heads today* gets measure one-half, *heads today and tails tomorrow* gets measure one-quarter, etc.

The second purpose served by the gong is to provide evidence that distinguishes *today* from every other day, e.g. *the gong rang today* or *the gong rang the day before yesterday*. However, one may endow *today* with such evidence by postulating and conditioning on irrelevant, almost sure position-specifying evidence such as the complete trajectory (relative to *today*) of a second, independent Eternal Coin.[3] Almost surely this trajectory will not be periodic, in which case its landscape of tosses will look different from *today*'s perspective than it does from any other day's.

We can now draw out the differences between our position and Dorr's. Suppose one learns P. Imagine next two hypothetical completions of the Coin's trajectory. On the first, the Coin lands *heads* seventeen more times (starting *today*) and then recreates the binary expansion of π. On the second, the Coin immediately recreates the binary expansion of π (starting *today*). We believe that Dorr would say that these two evidential scenarios "agree about what the world is like as a whole".[4] We disagree; they discriminate between two different worlds. Concomitantly, we advocate for a "self-location first" chance function.

One way to characterize this disagreement might be to say that we treat *today* as a rigid designator. The device of a second Coin brings this to light; in the second scenario (but not the first), the Coin's mirroring of the binary expansion of π begins on the day at which the past of the second Coin is equal to the second Coin's actual past now. That is, it begins *today*.

This discussion does, however, suggest an attempt to establish that $Cr(H|N) = 1$ for some *heads*-dominated null set N of trajectories determined by behavior on coordinates other than that of *today*, namely by choosing an N that is a set of uncentered worlds in the coin-first sense, i.e. a shift-invariant set. The resulting credence functions $Cr(\cdot|N)$ tend to fail countable additivity, and in some cases it may be possible to prevent this by stipulating that $Cr(H|N) = 1$. To illustrate the sort of countable additivity failure we have in mind, let O be the null event *the Coin lands tails exactly once, ever.* (O is sensitive to behavior of the Coin *today*, but it

[3]Dorr contemplates eliminating the trouble he perceives that evidence distinguishing today from every other day makes for ONE by brutely stipulating "that you have no such distinguishing evidence about today". Such a stipulation would not affect our case because it isn't necessary to observe that the trajectory of the second Coin is α (a function from the integers into $\{0,1\}$) in order to conditionalize on the trajectory of the second Coin being α (any more than it is necessary to observe that P in order to conditionalize on P, e.g.), and to conditionalize on such a trajectory involves treating it as hypothetical evidence.

[4]To emphasize the oddness of a "coin first" approach, note that on this view, two sequences of toss outcomes describe the same world whenever one is a temporal shift of the other. Therefore, any measurable set of uncentered worlds corresponds to a shift-invariant set of trajectories. Since Bernoulli shifts are ergodic, any such set has extreme measure (i.e. measure zero or one) or fails to be measurable (in the usual sense).

will serve as an example.) Conditionalization on O raises the same trouble as the gong; it seems that, conditional on O, one ought to be indifferent between any two possible displacements of one's position from that of the unique *heads* outcome.[5]

This phenomenon looks to be indigenous to conditionalization on shift-invariant N. What we seek, then, is a shift invariant set of trajectories determined by the Coin's past alone. Examples include *the number of tails tosses in the past is finite* and *the density of tails tosses in the past is zero.* To pursue the latter example, let D be the assertion that

$$\lim_{n\to\infty} \frac{\#\text{ heads in past } n \text{ days}}{n} = \lim_{n\to\infty} f_n = 1.$$

Let γ be the centered proposition that a certain n-tuple of toss outcomes will be instantiated in the next n days' tosses, starting from today. Based on what reasons for HALF have been offered (independence plus Principal Principle, for example), its defenders will surely hold that $Cr(\gamma|D) = 2^{-n}$. Let γ_{-n}, meanwhile, be the centered proposition that the same n-tuple of toss outcomes be instantiated in the *previous* n days. By reasoning in the spirit of UPDATING, it seems that $Cr(\gamma|D) = Cr_{+n}(\gamma_{-n}|D)$ should hold, where Cr_{+n} is the credence function that one will have n days in the future. But by shift invariance of D, a strong case can be made that also $Cr(\gamma_{-n}|D) = Cr_{+n}(\gamma_{-n}|D)$.[6] What that means, though, is that conditioning on D yields the uniform distribution over the *previous* n tosses, for any n. Assuming countable additivity of $Cr(\cdot|D)$, one could now, in a few lines, reach the disastrous conclusion $Cr(\neg D|D) = 1$.

One can indeed avoid this disaster by stipulating that $Cr(H|D) = 1$, but at high cost. For even if one is willing to sacrifice independence of the past and future tails, it still seems wrong that one should put $Cr(H|D) = 1$ when D is a collection of pasts σ for which there is nothing to defeat the presumption that $Cr(H|\sigma) = \frac{1}{2}$. Moreover (since D is shift invariant and assuming $Cr(H|D) = 1$) countable additivity has the unwelcome consequence that one ought to assign full credence, conditional on D, to *every* toss (past, present and future) landing *heads*. Given the Coin's fairness, that seems excessive.

For all this new take on the argument from evolving credences says, then, it may well remain a live option to maintain $Cr(H|D) = \frac{1}{2}$ and learn to live with mere finite additivity of $Cr(\cdot|D)$. Takers may find comfort in that $Cr(\cdot|D)$ will still be countably additive over events generated by tosses later than n days ago, for any n; presumably,

[5]Conditionalization on P doesn't have this effect. Expressed in wildcard notation, P corresponds to the set of sequences of the form $\ldots HHHH * * * * \ldots$, where the first wildcard $*$ occurs in the "today" position. So upon contemplating P, there isn't any "new point of reference" that you become lost with respect to. You know exactly where you are relative to P's features, namely at the day of the first toss that P fails to constrain.

[6]The reader will note a similarity to CONSTANCY, but also an important difference: unlike P, D is shift invariant and so not less likely to hold at the later time.

those are the tosses we care about. If that's right then any challenge to HALF must come from elsewhere.[7]

2. Appendix: Arbitrary conditional credences?

In a well researched paper (2015) serving in part as a reply to Dorr, Wayne C. Myrvold gives several fascinating examples of seemingly natural intuitions about conditionalization on measure zero sets coming into conflict with one another. But he also writes of the Eternal coin that:

> $P(H|\mathcal{P})(u) = P(H) = \frac{1}{2}$ for almost all u in our event space. But this doesn't preclude Dorr, or anyone else so inclined, from assigning the value 1, or any other value, to the probability of H conditional on the proposition P, or on any set of propositions comprising a set of measure zero. Distinct choices of this sort yield the same probabilities for all propositions about histories.

Here P is subjective probability (which we have been denoting Cr); the "event space" is the space of all doubly infinite $\{heads, tails\}$-valued sequences. (We again take the domain of these sequences to be the set of days, past present and future.) u ranges over one-sided sequences (possible pasts), i.e. the atoms of $\mathcal{P}$, the σ-algebra generated by the propositions

p_n: the coin landed *heads* n days ago.

Myrvold's contention is that the only constraints on $Cr(H|\mathcal{P})(u)$[8] are collective constraints borne by the full set of possible pasts u. Specifically, his position appears to be that these are constrained by (and only by) the requirement that

$$Cr(H \cap A) = \int_A Cr(H|\mathcal{P})(u)\, du \text{ for all } \mathcal{P}\text{-measurable sets } A. \tag{1}$$

If Myrvold is right, then one cannot argue at all for $Cr(H|\mathcal{P})(u) = \frac{1}{2}$ for any *single* past.

We suggest, against this, that (assuming there is no explicit stipulation about null event behavior, e.g. "the coin is such that if it has never landed *heads* in the past then it will never land *heads*"), probabilities of the sort $Cr(H \cap A)$ may in some cases constrain even single values $Cr(H|\mathcal{P})(u)$, albeit weakly. Specifically:

[7] One possible line of attack is to run a version of the Borel-Kolmogorov paradox (in which it is shown that, starting with the usual measure on the sphere, conditionalization on a great circle is ambiguous; see e.g. Mryvold 2015) in sequence space. We give a brief sketch. First identify (modulo null sets) the sequence space $\{0,1\}^{\mathbf{Z}}$ with the unit square, the negative and non-negative coordinates being viewed as binary expansions of x and y respectively. P may now be identified with the Eastern border of the square, and H with its upper half. Conditional on $\{(x,y) : x \geq y\}$, H has measure $\frac{1}{4}$. Indeed, H has measure $\frac{1}{4}$ conditional on any triangle with one vertex at $(1,1)$ and facing side on the x-axis. So, there is reason to assume that H has measure $\frac{1}{4}$ conditional on any line segment connecting $(1,1)$ and the x-axis. But the Eastern border, i.e. P, is such a line segment. We aren't sure that's enough ambiguity to overrule the Principal Principle, though.

[8] The notation Dorr employs, namely $Cr(H|u)$, is technically ambiguous in that no σ-algebra is specified. (See Myrvold (2015) for details.) Although Dorr surely had $\mathcal{P}$ in mind, only our response to Myrvold requires the more precise notation.

Convexity: Let u be a zero measure atom of $\mathcal{P}$. To maintain that $Cr(K|\mathcal{P})(u) = y \in \mathbf{R}$ on the basis of constraints on Cr alone, there must exist, for every neighborhood N of y, a $\mathcal{P}$-measurable set A such that $Cr(K|A) \in N$.

To be clear, *Convexity* is intended as a supplement to the collective constraints (1). (It is not a replacement!) At any rate...what can one say on behalf of *Convexity*?

Naively, $Cr(K|\mathcal{P})(u) = y$ means that, were you to learn that u is the case, your posterior in K would come to be y. Note, however, that there are no stipulations regarding null event behavior (countably many such would not constrain Cr), and u cannot be learned in finite time. In the coin toss example, this would involve checking the results of countably many tosses; such a case is typical.

Where one can learn u only "by degrees", it is natural that $Cr(K|\mathcal{P})(u) = y$ should at least be *consistent* with the aforementioned naive sentiment. In other words, it should imply that, were you to learn (by an infinite sequence of $\mathcal{P}$-measurable observations) that u is the case, your posterior in K *might* come to be (i.e. approach) y. Failing this condition, we see no way to ascribe to "$Cr(K|\mathcal{P})(u) = y$" any substantive meaning.

It is often assumed that the function $Cr(K|\mathcal{P})(\cdot)$ is only defined almost everywhere. More precise would be to say that there are multiple ways to define it, any two of which are guaranteed to agree only almost everywhere. A typical construction: let (P_n) be a generating sequence of progressively finer finite $\mathcal{P}$-measurable partitions. Letting u_n be the member of P_n containing u, put

$$Cr(K|\mathcal{P})(u) = \lim_n Cr(K|u_n).$$

Such a definition assigns a unique value for every (not just almost every) u. (For some null set of u where the limit doesn't exist, this "value" may be "does not exist".)

Of course if, you having carried out the above construction, we repeat it with a different generating sequence of partitions (Q_n), we may obtain $Cr(K|\mathcal{P})(u) = y_2$ where you had $Cr(K|\mathcal{P})(u) = y_1 \neq y_2$ (for some particular u or for some null set of u). Even so, each of us would still mean "if one were to learn (asymptotically, by means of our respective partition sequences) that u is the case, one's posterior in K would come to be (in the limit) y", and both definitions would satisfy *Convexity.*

In the case of the Eternal coin, $Cr(H|B) = \frac{1}{2}$ for every positive measure, $\mathcal{P}$ measurable set B, so *Convexity* requires $Cr(H|\mathcal{P})(u) = \frac{1}{2}$ for *every* (not just almost every) past u–which, we propose, is just what everyone knew all along.

References

Builes, David. (2019). A Paradox of Evidential Equivalence. *Mind*, forthcoming.

Dorr, C. (2010). The Eternal Coin: A puzzle about self-locating conditional credence. *Philosophical Perspectives* 24:189205.

Myrvold, Wayne C. (2015). You can't always get what you want: some considerations regarding conditional probabilities. *Erkenntnis* 80(3):573-603.

RMCCTCHN@MEMPHIS.EDU