SOLVING A PARADOX OF EVIDENTIAL EQUIVALENCE

CIAN DORR, JOHN HAWTHORNE, AND YOAAV ISAACS

Abstract. David Builes presents a paradox concerning how confident you should be that any given member of an infinite collection of fair coins landed heads, conditional on the information that they were all flipped and only finitely many of them landed heads. We argue that if you should have any conditional credence at all, it should be $\frac{1}{2}$.

1. Introduction

David Builes (2020) presents a paradox. When you know that a countable infinity of fair coins have been flipped and that only finitely many of them landed Heads, what should your credence be that some particular coin among them landed Heads? There are apparently strong reasons to believe that your credence should be 0 and there are apparently strong reasons to believe that your credence should be $\frac{1}{2}$.

We will argue that if you should have any credence at all that some particular coin landed Heads conditional on only finitely many coins having landed Heads, it should be $\frac{1}{2}$. We undermine the arguments for 0 that Builes presents by pointing out that for each of them there is a parallel and no less compelling argument for $\frac{1}{2}$. These arguments appeal to modes of reasoning which proponents of nontrivial credences conditional on credence-zero propositions have good reason to reject in any case. But there are other strong arguments for $\frac{1}{2}$ which do not appeal to these modes of reasoning, and do not correspond to any arguments for 0. We present two such arguments.

2. A Pair of Cases

Infinite cases are sometimes problematic in ways that their finite analogues are not. Here's a striking pair of cases.

Six coins: You are in a room with five other people. Each of you flips a coin without looking at the result. You know that all of the coin flips are fair and independent. You ask a nearby Oracle whether more coins landed tails than landed heads. The Oracle replies affirmatively. What should your credence be that your coin landed Heads?

This case is perfectly well-behaved. Initially, there are 64 possible outcomes—32 in which your coin lands heads and 32 in which it lands tails—and you should treat all as equiprobable. Of the 32 where your coin lands heads, one has five tails and five have four tails. Of the 32 where your coin lands tails, one has six tails, five have five tails, and ten have four tails. Given the Oracle's revelation, those 22 are the possibilities that remain. So your coin lands heads in six of the 22 live possibilities, and thus your credence that your coin landed heads should be $\frac{6}{22}$.

Now consider an infinitary analogue.

Infinitely Many Coins: You are in a room with countably infinitely many people. (Most of them are very tiny.) Each of you flips a coin without looking at the result. You know that all of the coin flips are fair and independent. You ask a nearby Oracle whether more coins landed tails than landed heads. The Oracle replies affirmatively. What should your credence be that your coin landed Heads?

This case is not perfectly well-behaved. Initially, there are uncountably infinitely many possible outcomes. Of those, countably infinitely many are outcomes in which only finitely many coins land heads. Given the Oracle's revelation, those are the possibilities that remain. Among them, countably infinitely many are possibilities in which your coin lands heads and countably infinitely many are possibilities in which your coin does not land heads. Combinatorics worked perfectly in the finite case, but it provides no guidance at all in the infinite case.

3. Builes' Paradox

It is not philosophically interesting that possibility-counting works as a guide to credence in a finite case but not in an otherwise similar infinite case. Fortunately, possibility-counting is not the only method available for assigning credences in these cases, so it's worth looking at what other methods suggest. What is philosophically interesting is that other methods seem to give inconsistent verdicts.

Builes (2020) argues that a case like Infinitely Many Coins presents a paradox. Builes presents an argument that your credence that your coin landed Heads should be 0 and a contrary argument that your credence that your coin landed Heads should be $\frac{1}{2}$. We will begin by sketching these arguments informally.

Argument for credence 0*:* Learning that only finitely many of the coins landed heads should make you much less confident that your coin landed heads. There's nothing special about your coin that could justify being confident that it is one of the vanishingly rare coins that landed heads. Just as learning that more coins landed tails than landed heads should make you less confident that your coin landed heads in the finite case, it should do so in the infinite case. Indeed, since the proportion of coins landing heads is 0, your credence that your coin landed heads should be 0.

Argument for credence $\frac{1}{2}$*:* If you had instead learned that only finitely many of the coins *other than yours* landed heads, that shouldn't affect your credences about your coin, since it wouldn't have anything to do with your coin. But you know for certain that the proposition that only finitely many of the coins other than yours landed heads is true if and only if the proposition that only finitely many of all the coins landed Heads is true, so they should have the same evidential import. Therefore learning that only finitely many of the coins landed heads shouldn't affect your credences about your coin. Since your initial credence that your coin would land heads was $\frac{1}{2}$, your credence that your coin landed heads should be $\frac{1}{2}$.

These each seem like good arguments. But they can't both be right.

Note that neither argument depends on your picking out the coin in question as *yours*. Both remain prima facie compelling if we assume that you know of some eternal, qualitative property that distinguishes your coin from the rest. Following Builes's suggestion, we will assume that this is the case, so as to set aside complications relating to the epistemology of self-locating and *de re* belief. For concreteness, let's say that your coin is a nickel and all the other coins are quarters.[1] It will also occasionally be convenient to assume that each coin is inscribed with a unique natural number: the nickel has number 0, and every positive number is inscribed (in increasingly tiny decimal notation) on

[1]In Builes's version, your coin is instead picked out as the only one tossed by a person wearing a red shirt. None of the arguments we will consider suggest that it could matter what the distinguishing qualitative property is, so long at there is no chance at the relevant time (before the coins are tossed) of its having a different extension than it actually has. But it is important that it be qualitative: if we attempt to run the puzzle using the proposition that x landed heads for some particular coin x, issues relating to the variety of 'guises' or 'modes of presentation' under which one could bear attitudes to such a proposition will loom large. See Hawthorne and Lasonen-Aarnio 2009 for difficulties which certain descriptive modes of presentation raise for standard principles about chance and credence, and Dorr 2010 for difficulties involving 'indexical' modes of presentation.

one of the quarters. And so that we can deal entirely in qualitative, eternal propositions, let's furthermore assume that these coins are all the coins there will ever be, and that each will be tossed exactly once.[2]

Finally we will dispense with the Oracle, and understand the puzzle as one about what your *conditional* credences should be, under circumstances where you have no relevant evidence beyond 'the setup': the proposition that is exactly one nickel and a countable infinity of quarters, each of which is tossed exactly once, and all these tosses are fair and independent. Following Builes, we can articulate three inconsistent claims, the second and third of which combine in the above argument for credence $\frac{1}{2}$.

Relevance: Your credence that the nickel lands heads conditional on the setup obtaining and at most finitely coins landing heads should be 0.

Independence: Your credence that the nickel lands heads conditional on the setup obtaining and at most finitely many quarters land heads should be $\frac{1}{2}$.

Equivalence: Your credence that the nickel lands heads conditional on the setup obtaining and at most finitely many coins

[2]In the version of the case where your only way of picking out your coin is by its relation to you (because your evidence doesn't tell you of any qualitative property that you instantiate that infinitely many others don't also instantiate), there is a distinctive new argument that your credence should be 0, based on a popular principle of *self-locating indifference*. In one version, this principle says that when E is the strongest qualitative property for which your evidence entails that you are E, and H is a qualitative proposition that entails that the proportion of all Es who are F is x, then your credence that you are F conditional on H should equal x if it is well-defined (see Elga 2004 and Dorr 2010; see also Manley MS and Dorr and Arntzenius 2017 for more powerful principles about *prior* credences from which this follows). When your evidence is just that you have a certain qualitative property that infinitely many of the other coin-owners also have, the proposition that finitely many coins landed heads entails that the proportion of people whose coins landed heads among those that have that property is 0; so self-locating indifference entails that your credence that you are such a person should be 0. (Dorr (2010) defends an analogous view about a very similar case.)

Friends of self-locating indifference could go on to argue that 0 must also be the right credence in the original version of the case, by appealing to the premise that it is irrelevant whether the description picking out the target coin involves a qualitative property or a relation to the agent. But this argument is dialectically weak. Since there is no prospect of a deduction of self-locating indifference from some consistent indifference principle that also applies to purely qualitative propositions, proponents of self-locating indifference had better think that self-locating credences are subject to distinctive rational constraints. And if there are such constraints, there should be no presumption that the two versions of the case are on a par.

> landing heads should be equal to your credence that the nickel lands heads conditional on the setup obtaining and at most finitely quarters landing heads.

Relevance, Independence, and Equivalence are each intuitively plausible. But they cannot all be true.[3] Which should be rejected?

4. The Easy Way Out

We are open to taking the easy way out and rejecting all three principles. All three principles presuppose that you should have credences conditional on certain credence-zero propositions. But it is not obvious that it is even *possible* for someone to have any credence in q conditional on p while having credence zero in p.[4] The difference between two credential states which differ only as regards credences conditional on credence-zero propositions seems to make no behavioural difference. The standard way of operationalising such differences connects facts about your credences conditional on some credence-zero p to facts about what your unconditional credences would be if you 'learned' p. But how *could* you learn p? Someone you trust could tell you that p—but unless you are crazy, your credence that they would never falsely tell you that p will be less than 1, so if you respond by

[3]If 'should' obeys a normal modal logic, the joint truth of all three principles would imply the absurd conclusion that 'You should do ϕ' is true for every ϕ. One might try to endorse all three principles by working with some weaker deontic logic designed to accommodate 'normative dilemmas'. But this way out of the puzzle does not seem promising enough to be worth exploring.

[4]We are speaking here and throughout of conditional credence *simpliciter*: a relation between a person, two propositions, and a number. We set aside concepts of conditional credence where it requires an additional argument, such as a partition or σ-algebra, such as those based on Kolmogorov's (1933) theory of 'regular conditional probability'. These relativised notions of conditional credence are not *quite* definable in terms of unconditional credence, but they come close: for details, see Easwaran 2019 (§2) or Myrvold 2015 (§4). The uncontroversial good standing of the relativized concepts means that rejecting non-relativised credences conditional on credence-zero propositions is a less radical option than it might initially seem: proponents of this option can interpret various claims that seem to be about unrelativised conditional credences as involving a tacit partition argument. A proponent of this view might conjecture that Builes's paradox, like the widely-discussed 'Borel-Kolmogorov' paradox (Myrvold 2015, §2), arises because *finitely many coins landed heads* belongs to two different salient partitions, relative to which your conditional credences should be 0 and 1/2 respectively. However, this conjecture is hard to substantiate, since it is entirely unclear what the two partitions could be.

conditionalizing on the fact that they told you p, you will just end up with with credence 1 that they spoke falsely.[5]

Only a behaviourist could regard such considerations as decisive objections to the possibility of having nontrivial credences conditional on credence-zero propositions.[6] But even non-behaviourists can think that there are deep constitutive links of some kind between credences and behavioural dispositions. The difficulty in operationalising psychological differences involving only credences conditional on credence-zero propositions may thus be taken as a warning that we don't really grasp what such differences amount to.[7] Paradoxes like Builes's arguably bolster this sceptical attitude: if certain putative psychological states seem to be subject to conflicting rational requirements, the explanation might be that the states didn't really make sense to begin with.

[5]Perhaps you could end up with credence 1 in p thanks to a carefully aimed blow to the back of the head, an appropriate course of brainwashing, or just a sudden leap of faith. And perhaps some possible cases like that would be correctly described as cases where you *learnt* (and thus came to know) p. Still, there is no saying what unconditional credences you would end up with if you learnt p thanks to some such odd process. You lack the stable dispositions that would constitute, or be explained by, your having particular credences conditional on p.

[6]Even behaviourists need not regard them as decisive, since they have the option of appealing to differences in language-related dispositions, e.g. dispositions to respond to questions of the form 'What is your conditional credence that ϕ given ψ?', or (as a referee suggested) to invitations like 'Suppose for the sake of argument that ψ', etc. But it is implausible that having nontrivial conditional credences requires having any familiarity with the relevant words in any language. And what conditional credences should we ascribe to someone whose dispositions vary depending on whether the questions are asked in English or French?

[7]The idea that there can't be differences in conditional credences without differences in unconditional credence doesn't have to go along with the idea that credences conditional on credence-zero propositions are undefined. One might instead say that all such credences are identical to 1. (In Popper's (1935) theory, the conditional probability function is well defined for all pairs of arguments, but there are some propositions, including contradictions, for which the function always gives the value 1 when they are the second argument.) Or, more interestingly, one might take the *naturalness* of certain partitions to ground the identification of one's credence in p conditional on a proposition q belonging to that partition with one's credence in p given q *relative to that partition*, using the partition-relative notion of conditional credence mentioned in note 4: in favourable cases, this can be uniquely fixed by unconditional credences by imposing a natural continuity requirement (see Easwaran 2019, §2.3.3). Or one might use some other method to reconstruct conditional credences from unconditional credences, e.g. looking at one's unconditional credences about conditional chances. But Builes's paradox will not be very gripping if we reconstruct conditional credences from unconditional credences in any of these ways, since there is no relevant controversy about what your unconditional credence function should look like.

Even if we set this kind of scepticism aside and grant that it would be *possible* for someone to assign a particular credence to *the nickel landed heads* conditional on *at most finitely many coins landed heads*, it is still not at all clear that it would be *rational* to do so. It is perfectly defensible to think that in some cases, although it would be possible to assign a credence to p conditional on q, it would be irrational to do so.[8] We can't think of any general principle that would mandate having a well-defined conditional credence in *the nickel landed heads* conditional on *at most finitely many coins landed heads*, except for principles that would make it rationally compulsory to assign a conditional credence to every pair of propositions one is capable of entertaining (perhaps with an exception for logical contradictions and propositions one treats in the same way one treats contradictions).[9] Assuming that it is possible for agents to have well-defined credences conditional on consistent propositions to which they assign credence zero, the view that they always *should* have such credences has certain attractions (see Elga 2010). But that view also leads to various kinds of unsettling arbitrariness in the assignment of conditional credences.[10] Many philosophers will think

[8]For example, it is arguably irrational to have any credences conditional on a contradiction—though if the contradiction is subtle, it may still be *possible*. It is widely, though not universally, held that there are some cases where it would be irrational to assign any *unconditional* credence to a certain proposition; if so, presumably the same phenomenon will arise with conditional credences. For an overview of such cases involving deliberation see Levi 2007. For an overview of such cases involving non-measurability see Hájek, Hawthorne, and Isaacs MS.

[9]Although Builes draws on Hájek (2003) to bolster his case that credence conditional on credence-zero propositions are sometimes well-defined, Hájek argues only that certain, especially natural credences conditional on credence-zero propositions are well-defined. The credences Builes' case relies on are not among those.

[10]Consider the collection C of all (finite or infinite) Boolean combinations of propositions of the form *coin n lands heads*. Any permutation π of the natural numbers generates a permutation π^* of C in an obvious way. It is tempting to think that your conditional credences as regards these should be *permutation invariant*, in the following sense: for each permutation π and any p and q in C, $Cr(p|q) = Cr(\pi^*(p)|\pi^*(q))$ if $Cr(p|q)$ is well-defined. But if $Cr(p|q)$ exists for every consistent p and q in C, there is no way to satisfy permutation invariance. Suppose for contradiction that you did. Let h be *exactly one coin landed heads*, or more strictly, the disjunction of all conjunctions which contain *coin n landed heads* for exactly one n and *coin m did not land heads* for every $m \neq n$. $\pi^*(h) = h$ for every permutation π. Since there is a permutation that interchanges *an even-numbered coin landed heads* and *an odd-numbered coin landed heads*, permutation invariance requires your credences in these two propositions conditional on h to be equal, hence $\frac{1}{2}$. But since there is also a permutation that maps *an even-numbered coin landed heads* to *an odd multiple of 3 landed heads* and maps this to *an odd non-multiple of 3 landed heads*, permutation invariance also requires your credences in these three

it more rational to avoid such arbitrariness by refraining from having well-defined conditional credences for some pairs of propositions. And once we allow that this is the right strategy in some cases, it would be quite natural to think that the solution to Builes's paradox is that the relevant conditional credences should not be well-defined, if that the arguments for Relevance, Independence, and Equivalence all looked strong on the assumption that they are well-defined.

But we are not convinced that this is the right way out. From now on, we will proceed on the assumption that you should have should have well-defined conditional credences for all the relevant pairs of propositions. We will also grant the stronger assumption that there are specific conditional credences that you should have. Working under these assumptions, we will argue that whereas Equivalence and Independence are supported by strong arguments, the arguments for Relevance merely dramatise some surprising structural features which proponents of credences conditional on credence-zero propositions already need to learn to live with. Thus, the paradox does little to advance the case against the possibility of rationally having well-defined conditional credences for the relevant propositions.

5. Equivalence: A Hill to Die On

As Builes explains at the end of his paper, on the assumption that the relevant conditional credences are well-defined, Equivalence can be derived straightforwardly from the Multiplicative Axiom for conditional probability, which says that when Cr is your conditional credence function,

$$Cr(p \wedge q|r) = Cr(p|q \wedge r)Cr(q|r)$$

whenever all three terms are defined. An immediate consequence of this principle is that when $Cr(q|r)$ and $Cr(r|q)$ are both 1,

$$Cr(p \wedge q|r) = Cr(p|q \wedge r) = Cr(p \wedge r|q)$$

But when $Cr(q|r) = Cr(r|q) = 1$, we also have $Cr(p|r) = Cr(p \wedge q|r)$ and $Cr(p|q) = Cr(p \wedge r|q)$.[11] So the Multiplicative Axiom has the important consequence that when your credence in each of two propositions conditional on the other is 1, they are interchangeable in the second argument of your conditional credence function:

$$Cr(p|q) = Cr(p|r) \text{ whenever } Cr(q|r) = Cr(r|q) = 1$$

propositions conditional on h to be equal, hence $\frac{1}{3}$. Thus you can obey permutation invariance only by lacking well-defined conditional credences for some of the pairs of propositions just considered.

[11]Since $Cr(p|r) = Cr(p \wedge q|r) + Cr(p \wedge \neg q|r)$ and $Cr(p \wedge \neg q|r) \leq Cr(\neg q|r) = 0$.

In the case at hand, the conjunction of the setup with *finitely many coins land heads* and the conjunction of the setup with *finitely many quarters land heads* logically entail one another, so your credence in each conditional on the other should be 1. This implies given the above theorem that your credences in *the nickel landed heads* conditional on these two propositions should be equal.[12]

The Multiplicative Axiom is intimately related to the principle that whenever your unconditional credences $Cr(p \wedge q)$ and $Cr(q)$ and your conditional credence $Cr(p|q)$ are all well-defined,

$$Cr(p \wedge q) = Cr(p|q)Cr(q)$$

Given that $Cr(p \wedge q) = 0$ whenever $Cr(q) = 0$, this is equivalent to the familiar 'ratio formula', according to which

$$Cr(p|q) = \frac{Cr(p \wedge q)}{Cr(q)}$$

whenever all three terms are defined and $Cr(q) > 0$. This formula plays an absolutely pervasive role in reasoning about probability; denying it would throw the entirety of Bayesian epistemology into disorder. But there seems to be something unstable about a view that accepts the ratio formula (as a constraint on rational conditional and unconditional credence) while rejecting the Multiplicative Axiom. If one's credences conditional on r correspond in the usual way to one's dispositions to respond to 'learning' r, and one violates the Multiplicative Axiom with respect to p, q, and r, then one will be *disposed* to violate the ratio formula upon learning r. If this is a disposition to do something irrational, it seems like a departure from rationality in its own right.

The Multiplicative Axiom seems moreover to do important work in pinning down the very concept of conditional credence. It provides a source of discipline that those who reject the project of defining conditional credence in terms of unconditional credence can ill afford to let go of. Without the Multiplicative Axiom, there is so little left to constrain our theorising about rational credences conditional on credence-zero propositions that the entire topic risks seeming like a dead end.

6. Two Arguments for Independence

We will give two arguments for Independence. The first appeals only to the symmetries of the setup, and is novel; the second deals

[12]The fact that these propositions logically entail one another means that on many moderately coarse-grained theories of propositions, they are one and the same, in which case we can derive Equivalence directly from Leibniz's Law, without any need to appeal to the Multiplicative Axiom.

with objective chance, and is closely related to arguments which Builes considers.

(i) Symmetry. Suppose we apply some paint to the coins. We paint the Heads side of the quarters green and the Tails side of the quarters blue, and we paint the Heads side of the nickel blue and the Tails side of the nickel green. Nothing in this modification of the setup privileges the property of landing heads over the property of landing green. So, if there is a particular number x which should be your credence that the nickel landed heads conditional on the conjunction of the painted-coin setup with *finitely many quarters landed heads*, then x should also be your credence that the nickel landed *green* conditional on the conjunction of the painted-coin setup with *finitely many quarters landed green.* Since those two conjunctions are logically equivalent, the Multiplicative Axiom entails that they are interchangeable in the second argument of your conditional credence function, as explained in §5. So, x should also be your credence in *the nickel landed green* conditional on the conjunction of the painted-coin setup with *finitely many quarters landed heads*. But since the painted-coin setup specifies how the nickel is painted, you should have credence 0 conditional on this conjunction that the nickel both landed heads and landed green, and credence 1 that it either landed heads or landed green. So the only way to assign the same conditional credence to *the nickel landed heads* and *the nickel landed green* is to assign $\frac{1}{2}$ to both. And if your credence should be $\frac{1}{2}$ in the painted-coin variant of the setup, it should surely also be $\frac{1}{2}$ in the original version.[13]

One could resist this argument by insisting on some interpretation of 'the coin flips are fair and independent' on which it is spelled out in terms of 'heads' and 'tails' and thus might be thought to constrain your credences about heads and tails in a certain way that it does not constrain your credences about blue and green. But this is unpromising. Saying that a coin flip is 'fair' just means that the two sides of the coin involved have the same chance of coming up (conditional on its being tossed): it has nothing to do with what kinds of decorations, if any, distinguish the two sides. And whatever exactly 'independent' means, it is presumably some general relation that can obtain between chancy processes of all sorts, including processes that do not involve coins at all.

---

[13]In fact the paint isn't needed: we can run the symmetry argument directly in the original case by replacing *landing green* with *either landing heads and being a quarter or landing tails and being a nickel*.

One could alternatively resist the argument by denying that there is any particular number x that should be your credence that the nickel landed heads conditional on the setup obtaining and finitely many quarters having landed heads. Perhaps there are many different conditional credences which you could permissibly assign to this pair of propositions. Or perhaps which particular credence you should have depends on details of your situation which were left open by our stipulations, such as your past experience with similar coins. However, this way of resisting the argument is not open to proponents of the combination of Relevance and Equivalence, who are committed to the view that there is a particular credence you should have, namely 0.

(ii) Chance. If we allow for rational people to have well-defined credences conditional on propositions in which their unconditional credence is zero, it is plausible that we should also allow for propositions to have well-defined *chances* conditional on propositions whose unconditional chance is zero. And if we allow for primitive chances conditional on chance-zero propositions, it is most natural to spell out 'the coin flips are fair and independent' as follows: for any coin x, any side x' of x, and any consistent finite or infinite Boolean combination p of propositions each of which is of the form y *on* y' for some coin y distinct from x and some side y' of y:

$$\text{Chance}_t(x \textit{ lands on } x'|p) = \tfrac{1}{2} \tag{*}$$

where t is a time before any coins have been tossed, but after it has been determined that all of them will be tossed exactly once.[14]

So in particular, the setup entails that (*) holds when x is the nickel, x' is its 'tails' side, and p is the disjunction of all conjunctions of collections of propositions which include exactly one proposition of the form y *lands on* y' for each quarter y, where y' is the 'tails' side for at most finitely many of the quarters and otherwise is the 'heads' side. But there is no chance at the relevant time of anything changing whether it is, a quarter, a nickel, or not a coin at all, or of any side of a coin changing whether it displays heads or tails. So the setup also entails

[14]'Fair and independent' could also be given a weaker interpretation where (*) need only hold when p is a consistent *finite* Boolean combination of propositions about coins other than x. But the stronger interpretation is more natural, and using the weaker one would get us bogged down in an unilluminating debate about how confident you should be that the coin flips are independent in the strong sense, conditional on their being fair and independent in the weak sense.

that:

$$\begin{aligned}
1 &= \text{Chance}_t(p \mid \textit{finitely many quarters land heads}) \\
&= \text{Chance}_t(\textit{finitely many quarters land heads} \mid p) \\
&= \text{Chance}_t(\textit{the nickel lands heads} \mid p \textit{ and } x \textit{ lands on } x') \\
&= \text{Chance}_t(x \textit{ lands on } x' \mid p \textit{ and the nickel lands heads})
\end{aligned}$$

Applying the Multiplicative Axiom for conditional chance, we can deduce that

$$\text{Chance}_t(\textit{the nickel lands heads} \mid \textit{finitely many quarters land heads}) = \tfrac{1}{2}$$

But you have no relevant evidence about how the coins landed beyond the setup. Given the Principal Principle connecting chance with rational credence (Lewis 1980)—or, more precisely, the natural extension of this principle to *conditional* chance and *conditional* credence—it follows that whenever the setup entails that the chance at some time of p conditional on q is x, your credence in p conditional on the conjunction of q with the setup should also be x.[15] In particular, your credence in *the nickel lands heads* conditional on the conjunction of the setup with *finitely many quarters land heads* should be 1/2, in accordance with Independence.[16] Of course, the Principal Principle is not completely uncontroversial, even for qualitative propositions. But the main objections to it in the literature turn on a kind of 'undermining'

[15]The relevant version of the unconditional Principal Principle says that where Cr is any rational prior credence function, p is any proposition, and h is any proposition entirely about history up to t and the chances at and before t that entails that $\text{Chance}_t(p) = x$, $C(p|h) = x$ if it is well-defined. The conditional chance version says that if p and q are any two propositions and h is any proposition entirely about history up to t and the *conditional* chances at and before t that entails that $\text{Chance}_t(p|q) = x$, $C(p|q \wedge h) = x$ if it is well-defined. In our application, h is the setup. Since you have no relevant evidence, we can treat the question what your credence should be as tantamount to the question what a rational prior credence would be.

[16]One could run a parallel argument in the self-locating version of the case discussed in note 2, replacing *nickel* with *coin owned by you*. The conclusion of that argument conflicts with the credence of 0 prescribed by self-locating indifference. But it is not news that self-locating indifference is inconsistent with the unrestricted conditional-chance Principal Principle: the central example in Dorr (2010) also shows this. Even for those firmly attached to self-locating indifference, rejecting the conditional-chance Principal Principle entirely seems like an overreaction to this conflict, since it leaves the concept of primitive conditional chance looking worryingly underconstrained. If we don't want to give up on primitive conditional chance and credence, it seems more reasonable to hold on to the restriction of the Principal Principle to qualitative eternal propositions, which is consistent with self-locating indifference, and is all we need to argue for Independence.

that can happen in Humean theories of chance, which does not look like a promising line of resistance for proponents of Relevance.[17]

Since the chance-based argument for Independence relies crucially on the setup's entailing that there is no chance of objects changing what kind of coin they are, etc., it is natural to wonder how things go if this assumption is false. The answer, interestingly, is that in some cases like this, the Principal Principle actually requires that your credence that the nickel landed heads conditional on finitely many quarters landing heads should be *zero*, if it is well-defined. Suppose that all coins start out as quarters, and are tossed between t_0 and t_1. Between t_1 and t_2, a demon will randomly choose exactly one coin to transform into a nickel, in such a way that each coin has chance zero at t_1 of ending up as a nickel.[18] Since chances are necessarily finitely additive, this demon-setup entails:

(i) Every finite set of coins has chance 0 at t_1 of containing a coin that will become a nickel.

Moreover, since truths about which coins did and did not land heads all have chance 1 at t_1, and truths about set-membership have chance 1 at all times, the demon-setup also entails:

(ii) Whichever set S contains all and only the coins that landed heads is such that the chance at t_1 that a member of S becomes a nickel equals the chance at t_1 that a coin that landed heads becomes a nickel.

---

[17]See Lewis 1994 for the problem of undermining and a revision of the Principal Principle that avoids it. Note that Humean theories of chance raise a prima facie worry about the very consistency of the setup with the proposition that finitely many quarters landed Heads, since they suggest that an infinite collection of coin flips only finitely many of which landed heads might ipso facto not be fair and independent. And even if we admit the possibility of the case, Humeanism does nothing to suggest that your conditional credence should be *lower* than $\frac{1}{2}$—if anything, it might motivate making it higher, on the grounds that having any one coin land heads makes it just a little easier for the overall Humean mosaic to be one of those that makes the coins fair and independent.

[18]This example requires the possible failure of the countable additivity principle for unconditional chance. For a version compatible with the necessity of this principle, we could imagine that there are *continuum* many coins, one for each real number in the unit interval, so that the demon can choose one by throwing a very pointy dart at a target and measuring the distance to the centre. Alternatively, we could just have the demon toss each coin a second time and transform all those that land heads into nickels, and replace all our references to unconditional chance with references to chances conditional on the proposition that only one coin lands heads when tossed by the demon.

The conjunction of (i) and (ii) with the proposition that the set of coins that land heads is finite entails

(iii) The chance at t_1 that a coin that landed heads becomes a nickel is 0.

You should therefore have credence 1 in (iii) conditional on the conjunction of the demon-setup with *finitely many coins landed heads*. So, applying the Principal Principle to your credences about the chances at t_1, it follows that your credence that a coin that landed heads became a nickel, conditional on that conjunction, should also be zero: Relevance is true for this variant of the case.[19]

By granting that your conditional credence should be zero in the demon version of the case, proponents of Independence can arguably undercut some of the intuitive appeal of Relevance: this appeal may be partly based on mistakenly assuming that the factors that distinguish the two versions don't make any difference.[20] Still, the idea that the two versions require different credences is already quite surprising, especially when we notice that the difference is entirely a matter of temporal order: if, instead, we have the demon choose a coin between t_0 and t_1 and have the coin flips occur between t_1 and t_2, the argument from the Principal Principle to Independence will go through just as before (taking the relevant time t to be t_1). That is strange, since either way there need be no *causal* connection between the demon's choice-process and the coin-tossing processes.[21] This *could* be regarded as providing further support for the claim that you should not have a well-defined conditional credence either way. However, once we reflect on the central role of time in the operative concept of chance, it becomes less surprising to find that questions of temporal ordering make

[19]Our earlier symmetry-based argument for Independence does not conflict with this result. If the heads sides of the coins are initially blue and the tails sides green, but the demon will swap the colours on whichever coin he turns into a nickel, there is a relevant chance-theoretic asymmetry between *heads* and *green*: the analogue of (ii) substituting 'landed on the side that would end up green' for 'landed heads' is false.

[20]The conclusion may also be reassuring for proponents of Independence who accept self-locating indifference and are thus committed to Relevance in the self-locating version of the case (see note 2). If the self-locating version is analogous to one of the two cases, it is not so surprising that it should be the demon version rather than the original version, especially since the idea of self-locating indifference is often heuristically expressed by formulae like 'You should treat yourself as a random sample' (Bostrom 2002, Manley MS).

[21]To play this up, we could replace the demon with a random number-choosing device located far from the coins, while replacing 'nickel' with 'coin whose number matches the one chosen by the device'.

a difference. Even setting questions about credences aside, a plausible 'reflection principle' concerning the relation between chances at earlier times and chances at later times entails that when the coins are tossed first, the chance at t_0 of *a nickel will have landed heads* conditional on *finitely many coins land heads* is 0, whereas in the version where the demon goes first, it is $\frac{1}{2}$.[22] Given this striking chance-theoretic contrast between the two versions of the case, the claim that there is a corresponding rationality-theoretic contrast looks much less objectionable.[23]

7. Is there a case for Relevance?

It thus remains for us to address the arguments for Relevance developed in Builes's paper. We will show that these arguments provide no *differential* support to Relevance over Independence, since exactly parallel arguments can be given for Independence. The arguments just bring out surprising structural features which are inevitable consequences of having *any* conditional credence function which assigns credences conditional on some credence-zero propositions and meets some minimal further constraints. Once this is recognised, the arguments have no force either way. By contrast, there are no arguments for Relevance analogous to our earlier arguments for Independence. There is thus no parity between the two principles; Independence is in strictly better shape than Relevance.

Builes gives three arguments for Relevance: the first turns on considerations related to accuracy; the second involves a Dutch book; and the third turns on general principles about the relationship between conditional and unconditional probabilities. We will discuss them in reverse order.

(i) Conglomerability. Assuming that you should have any credences conditional on any credence-zero propositions, it seems plausible that

[22] The reflection principle we have in mind says that if h is entirely about history up to t^+ and chances at and before t^+, and entails $\text{Chance}_{t^+}(p|q) = x$, then $\text{Chance}_{t^-}(p|q \wedge h) = x$ if it is well-defined. Since this principle has the same form as our conditional-chance Principal Principle, we can use it to run analogues of our earlier arguments substituting *chance at* t_0 for *rational prior credence*.

[23] We might wonder what happens in a version of the case where the demon's choice and all the coin-tosses occur simultaneously. The answer is that in that case, our stipulations simply do not settle the question what the chance was that a nickel would land heads. We could consistently elaborate the setup so as to entail that this chance is 0, or $\frac{1}{2}$, or anything in between. And the question what our credence should be conditional on the un-elaborated setup will depend on how we should divide our credence among these different elaborations. (Cf. the discussion in Dorr 2010 (§6) of a structurally analogous case involving two different chancy processes, namely the tossing of a single coin infinitely many times and the ringing of a gong.)

for each natural number n, your credence that the nickel landed heads conditional on the proposition that exactly n coins landed heads should be 0. For as Builes argues, it would be bizarre for your credence in *the nickel landed heads* conditional on *exactly n coins landed heads* to be greater in the case where there are infinitely many coins in total than it would be if the total number of coins was some finite number $m \geq n$; but since your credence should uncontroversially be n/m in the latter case, the only way to satisfy this constraint is for it to be 0 in the infinite case. But the proposition that only finitely many of the coins landed Heads is logically equivalent to the disjunction of all propositions of the form *exactly n coins landed Heads*. If your credence that the nickel lands heads should be 0 conditional on each of these propositions, mustn't it also be 0 conditional on their disjunction? Suppose you have already learnt that the disjunction is true but are still waiting to learn which disjunct is true. Shouldn't you set your credence to 0 now, given that you know you will do so later when you have learnt one of the disjuncts?

This argument is in effect an appeal to a very natural principle—*conditional conglomerability*. Conditional conglomerability says that given any propositions p and q and any collection X of pairwise inconsistent propositions whose disjunction is equivalent to q, if you are rational and your credence in p conditional on each member of X belongs to a certain closed interval, then your credence in p conditional on q also belongs to that closed interval.[24] In this case, p is *the nickel landed heads*; q is *at most finitely many coins landed heads*; X is the collection of all propositions of the form *exactly n coins landed heads*; and the closed interval is the singleton set $\{0\}$. Since your credence in p conditional on each member of X belongs to this interval, conditional conglomerability requires your credence in p conditional on q also to belong to it, i.e. to be 0. Conditional conglomerability thus dictates Relevance.

But it does not matter what conditional conglomerability dictates, because failures of conditional conglomerability are almost inevitable once we allow for primitive conditional credences on some zero-credence

[24]Conditional conglomerability is the conditional-probability generalisation of the concept of conglomerability introduced by de Finetti (1972, ch. 5): that is equivalent to the special case of conditional conglomerability where q is logically necessary, so that we can replace 'your credence in p conditional on q' with 'your credence in p'. If you violate conditional conglomerability with respect to p, q, and X, you will be disposed, upon learning q, to violate conglomerability with respect to a collection of propositions derived from X by disjoining the negation of q with an arbitrarily chosen member of X.

propositions. To illustrate why, let p_i for each natural number i be the proposition *the only coin to land heads was numbered* $2i$*,* $4i+1$*, or* $4i+3$, and let q_i be *the only coin to land heads was numbered* $2i+1$*,* $4i$*, or* $4i+2$. The disjunction of the p_i and the disjunction of the q_i are both equivalent to *exactly one coin landed heads*. Once we allow for primitive credences conditional on credence-zero propositions, it is hard to deny that your credence in *an odd-numbered coin landed heads* should be $\frac{2}{3}$ conditional on each p_i and should be $\frac{1}{3}$ conditional on each q_i.[25] But then conditional conglomerability imposes the unsatisfiable requirement that your credence in *an odd-numbered coin landed heads* conditional on *exactly one coin landed heads* be both $\frac{1}{3}$ and $\frac{2}{3}$.[26]

[25]This could be argued for in several ways. One strategy appeals to 'total outcome indifference': the plausible thesis that you should treat any two propositions that fully specify how each coin lands as equiprobable, in the strong sense of assigning them equal credence conditional on anything they both entail. This implies that conditional on p_i you should assign the same credence (namely $\frac{1}{3}$) to each of *only* $2i$ *landed heads*, *only* $4i+1$ *landed heads*, and *only* $4i+3$ *landed heads*; similarly for q_i. Another strategy appeals to the permutation invariance principle from note 10. Since the permutation that maps $2i$ to $4i+1$, $4i+1$ to $4i+3$, $4i+3$ to $2i$ and leaves every other number alone maps p_i to itself, invariance under this permutation requires your credences in *only* $2i$ *landed heads*, *only* $4i+1$ *landed heads*, and *only* $4i+3$ *landed heads* conditional on p_i to be equal; similarly for q_i.

[26]Schervish, Seidenfeld, and Kadane (1984) prove some general theorems which show how hard it would be for conditional conglomerability to hold once we allow some well-defined probabilities conditional on probability-zero propositions. One theorem says that unconditional conglomerability must fail (with respect to some countable partition) in any probability function for which countable additivity fails, so long the range of that probability function is an infinite subset of the unit interval. But whatever you want to say about countable additivity for *unconditional* credences, countable additivity for *conditional* credences seems extremely implausible (assuming we allow for well-defined, non-partition-relative credences conditional on zero-credence propositions). For example, if you thought that your credences conditional on *exactly one coin landed heads* should be countably additive, you would have to think that for some n, your credence in *only of the first n coins landed heads* conditional on *exactly one coins landed heads* should be greater than 0.99, which seems bizarre. And the idea that there are only finitely many numbers that are your credence in some proposition conditional on *exactly one coin landed heads* also seems completely implausible: for example, your credence in each proposition of the form *life evolved on exactly n planets* conditional on *exactly one coin landed heads* should plausibly be positive. Further theorems due to Schervish, Seidenfeld, and Kadane (1984, 2016) show that even if we were prepared to accept one of these very implausible things, maintaining conditional conglomerability in full generality would require severe restrictions on the set of pairs of propositions with well-defined conditional probabilities.

So, appeals to conditional conglomerability cannot be relied upon when we are dealing with credences conditional on credence-zero propositions. The only question is *where* and *how* it should fail. Builes is right that rejecting Relevance leads under natural assumptions to failures of conditional conglomerability with respect to the collection of propositions of the form *exactly n coins land heads*, whereas accepting Relevance allows us to maintain conditional conglomerability for this particular collection. But there are other collections for which rejecting Independence leads under equally natural assumptions to failures of conditional conglomerability, whereas accepting Independence allows us to maintain conditional conglomerability. One such collection comprises the propositions of the form *exactly n quarters land heads*. Your credence that the nickel lands heads conditional on each of these propositions should plausibly be $\frac{1}{2}$: the reasoning behind Relevance does nothing to undermine the plausibility of this. But the disjunction of all these propositions is equivalent to *only finitely many quarters land heads*. So to satisfy conditional conglomerability with respect to this collection, your credence in *the nickel landed heads* conditional on *only finitely many quarters landed heads* would have to be $\frac{1}{2}$, as Independence says it should be.[27] So, pending some new argument that conditional conglomerability with respect to propositions about how many *coins* landed heads is to be favoured over conditional conglomerability with respect to propositions about how many *quarters* landed heads, appeals to conditional conglomerability do nothing to favour Relevance over Independence. At most, appeals to conditional conglomerability could be used to support the view that the relevant conditional credences should be ill-defined (thus rejecting both Relevance and Independence).

(ii) Dutch Books. Builes shows that accepting Independence and rejecting Relevance disposes you to accept collections of bets which guarantee a loss. Suppose that after having learned that only finitely many of the coins landed heads, your credence that each particular numbered coin landed heads is still $\frac{1}{2}$. In that case, you should be willing to bet at favorable odds that each coin landed heads. But you are certain that any infinite collection of such bets contains infinitely many losing bets and only finitely many winning bets, thus guaranteeing an

[27]Another such collection contains, for each finite set S of positive numbers, the proposition—call it z_S—that all and only the quarters in S landed heads. The claim that your credence in *the nickel landed heads* conditional on each z_S should be $\frac{1}{2}$ follows *total outcome indifference* (see note 25). But the disjunction of the z_S is also equivalent to *finitely many quarters land heads*, so conditional conglomerability with respect to this collection also dictates Independence.

infinite loss. In contrast, if you follow Relevance, you will reject all the bets, and thus lose nothing.

This Dutch Book argument is, in effect, just highlighting a particular failure of conglomerability. Whenever conglomerability fails, an agent can fall prey to such an infinitary Dutch Book.[28] It should therefore be unsurprising, given that failures of conglomerability will arise upon learning that only finitely many coins landed heads whether we follow Relevance or Independence, to find that analogous Dutch Books can be made against agents whose credences conform to Relevance. Let p_n be *none of the first n coins lands heads*. Having learned that finitely many quarters landed heads, followers of Relevance will have credence 0 that each particular coin landed heads, and will therefore assign each p_n credence 1. But the conjunction of all the p_n is equivalent to *no coins landed heads*, and under plausible assumptions, your credence in this proposition should be 0 even after conditionalizing on *finitely many quarters landed heads*.[29] Now consider the collection of bets $b_1, b_2, \ldots$, where b_n pays out \$2 if p_n is true and costs \$1 if p_n is false. If you follow Relevance, you will accept all these bets, even though you are sure you will win only finitely many of those bets and lose infinitely many of them, thus losing an infinite amount of money. By contrast, if you follow Independence, your credences in $p_1, p_2, p_3 \ldots$ will be $\frac{1}{2}$, $\frac{1}{4}$, $\frac{1}{8} \ldots$, so you will refuse all the bets after b_1. In that case you are not guaranteed a loss: you will come out ahead if the first coin lands tails. Dutch Book arguments thus provide no support to Relevance over Independence.

(iii) Accuracy. Builes shows that if you follow Independence, there will be a collection of propositions for which you are guaranteed to have infinite total inaccuracy given standard strictly proper scoring rules, whereas if you follow Relevance you will have finite total inaccuracy arising from those propositions. Consider the collection of all propositions of the form *coin n lands heads*. Upon learning that only finitely

[28]Note that these Dutch books have a more respectable structure than those employed in McGee (1999). Unlike those, these Dutch books can be constructed with unconditional bets without infinite quantities of money being both gained and lost.

[29]This follows from the principle of total outcome indifference (note 25). Given the Multiplicative Axiom, that principle entails that conditional on *finitely many quarters landed heads*, your credence in *no quarters land heads* must be equal to your credence in *the quarter numbered n was the only one to land heads*, for each n. Since all these propositions are pairwise inconsistent, the only way this can happen is for all of them to be zero.

many coins landed heads, you know that finitely many of those propositions are true and infinitely many of those propositions are false. If you assign credence $\frac{1}{2}$ to each one—as you presumably will if you follow Independence—you are thus guaranteed to accumulate infinite inaccuracy (assuming that having credence $\frac{1}{2}$ in any false proposition makes the same positive contribution to total inaccuracy). By contrast if you assign credence 0 to each of those propositions, you will accumulate only a finite amount of inaccuracy (assuming that having credence 0 in a false proposition makes no contribution to inaccuracy).[30]

It is somewhat difficult to know what to make of this argument. When there are infinitely many propositions around, it is almost inevitable that you will have an infinite total inaccuracy score. It is not clear why picking some particular collection of propositions where one view leads to infinite total inaccuracy and the other view doesn't would have even *prima facie* force.

Even granting that this sort of accuracy argument has force, it too highlights a particular failure of conglomerability. Whenever conglomerability fails, an agent can fall prey to such an infinitary accuracy argument.[31] It should therefore be unsurprising, given that failures of conglomerability arise either way, to find that analogous accuracy arguments can be formulated against agents who follow Relevance. Let q_n be the proposition that at least one coin lands heads and none of the first n coins lands heads. Suppose that you accept Relevance and learn that at most finitely many coins landed heads. Presumably you will then have credence 0 that coin n landed heads for each n, but have credence 1 that at least one coin landed heads, and thus assign credence 1 to each q_n. But you are sure that at most finitely many of the q_n are true and the rest are false, so you are guaranteed to accumulate infinite inaccuracy with respect to these propositions, assuming that having credence 1 in any false proposition makes the same positive contribution to inaccuracy. In contrast, if you follow Independence you can (and presumably will) assign credence $(\frac{1}{2})^n$ to q_n. Under plausible assumptions about how inaccuracy is measured, this credence function will accumulate only a finite amount of total inaccuracy as far as the q_n are concerned: your positive credences in the infinitely many false q_ns sum to some finite amount, and hence the contributions of these

[30] The most commonly discussed strictly proper scoring rules satisfy both of these assumptions: see Joyce (1998, 2009).

[31] See Easwaran 2013 for arguments that accuracy-theoretic considerations require conglomerability.

credences to your inaccuracy score also sum to some finite amount.[32] Accuracy arguments thus provide no differential support to Relevance over Independence.

8. Conclusion

Builes' paradox is initially gripping because we are prone to make conditional conglomerability inferences. 'Finitely many quarters land heads' makes salient the 'How many quarters land heads?' partition and primes us to do conditional conglomerability reasoning for it. 'Finitely many coins land heads' instead makes salient the 'How many coins land heads?' partition and primes us to do conditional conglomerability reasoning for it. But we know conditional conglomerability cannot be good in general in the primitive conditional probability framework, so it is simply a mistake to presuppose it here. Once we overcome our temptation to make this mistake, it is straightforward to see that Relevance must be false. What remains to be seen is whether Independence is true, or rather trivially false because the relevant conditional credences should not be well-defined.[33]

Cian Dorr
New York University
cian.dorr@nyu.edu

John Hawthorne
University of Southern California and Australian Catholic University
johnhawthor@usc.edu

Yoaav Isaacs
Baylor University
Yoaav_Isaacs@baylor.edu

[32]It is sufficient for there to be some positive constants c and ϵ, such that whenever $x < \epsilon$, the inaccuracy contributed by having credence x in a false proposition is no greater than cx. For this it is in turn sufficient for the function specifying the inaccuracy contribution of a credence in any false proposition to be differentiable on some closed interval $[0, \epsilon]$: then c can be taken to be the maximum value of the differential of the function on this interval. The strictly proper scoring rules mentioned by Joyce (1998, 2009) all have this property.

[33]Thanks to an anonymous referee, whose questions and suggestions helped us to sharpen the discussions of variants of the original case in note 2 and §6.ii.

References

Arntzenius, Frank, Adam Elga, and John Hawthorne 2004, ‘Bayesianism, Infinite Decisions, and Binding’, *Mind* 113 (450), pp. 251–83.

Bostrom, Nick 2002, *Anthropic Bias: Observation Selection Effects in Science and Philosophy* (New York: Routledge).

Builes, David 2020, ‘A Paradox of Evidential Equivalence’, *Mind* 129 (513), pp. 113–27.

de Finetti, Bruno 1972, *Probability, Induction, and Statistics: The Art of Guessing* (New York: John Wiley).

Dorr, Cian 2010, ‘The Eternal Coin: A Puzzle about Self-locating Conditional Credence’, *Philosophical Perspectives* 24 (1), pp. 189–205.

—— and Frank Arntzenius 2017, ‘Self-Locating Priors and Cosmological Measures’, in Khalil Chamcham, John Barrow, Simon Saunders, and Joe Silk (eds.), *The Philosophy of Cosmology* (Cambridge: Cambridge University Press), pp. 396–428.

Easwaran, Kenny 2013, ‘Expected Accuracy Supports Conditionalization—and Conglomerability and Reflection’, *Philosophy of Science* 80 (1), pp. 119–142.

—— 2019, ‘Conditional Probabilities’, In Richard Pettigrew and Jonathan Weisberg (eds.), *Open Handbook of Formal Epistemology*, PhilPapers Foundation, pp. 131–98, <https://philpapers.org/rec/PETTOH-2>.

Elga, Adam 2004, ‘Defeating Dr. Evil with Self-Locating Belief’, *Philosophy and Phenomenological Research* 69 (2), pp. 383–96.

—— 2010, ‘Subjective Probabilities Should Be Sharp’, *Philosopher’s Imprint* 10.

Hájek, Alan 2003, ‘What Conditional Probability Could Not Be’, *Synthese* 137 (3), pp. 273–323.

——, John Hawthorne, and Yoaav Isaacs MS, ‘A Measure-Theoretic Vindication of Imprecise Credence’. Unpublished manuscript.

Hawthorne, John and Maria Lasonen-Aarnio 2009, ‘Knowledge and Objective Chance’, in Patrick Greenough and Duncan Pritchard (eds.), *Williamson on Knowledge* (Oxford: Oxford University Press), pp. 92–108.

Joyce, James M. 1998, ‘A Nonpragmatic Vindication of Probabilism’, *Philosophy of Science* 65 (4), pp. 575–603.

—— 2009, ‘Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief’, in Franz Huber and Christoph Schmidt-Petri (eds.), *Degrees of Belief*, Synthese Library (New York: Springer), pp. 263–97.

Kolmogorov, Andrei Nikolaevich 1933, *Grundbegriffe der Wahrscheinlichkeitrechnung* (Berlin: Springer), trans. Nathan Morrison as *Foundations of Probability* (New York: Chelsea Publishing Company), 1950.

Levi, Isaac 2007, 'Deliberation Does Crowd Out Prediction', in Toni Rønnow-Rasmussen, Björn Petersson, Jonas Josefsson, and Dan Egonsson (eds.) *Hommage à Wlodek: 60 Philosophical Papers Dedicated to Wlodek Rabinowicz* (Lund: Lund University), <www.fil.lu.se/hommageawlodek>.

Lewis, David K. 1980, 'A Subjectivist's Guide to Objective Chance', in Richard C. Jeffrey (ed.), *Studies in Inductive Logic and Probability, Volume II* (Berkeley: University of California Press), pp. 263–93.

—— 1994, 'Humean Supervenience Debugged', *Mind* 103 (412), pp. 473–90.

Manley, David MS, 'On Being a Random Sample'. Unpublished manuscript: <philpapers.org/rec/MANOBA>.

McGee, Vann 1999, 'An Airtight Dutch Book', *Analysis* 59 (4), pp. 257–65.

Myrvold, Wayne C. 2015, 'You Can't Always Get What You Want: Some Considerations Regarding Conditional Probabilities', *Erkenntnis* 80 (3), pp. 573–603.

Popper, Karl 1935, *Logik der Forschung. Zur Erkenntnistheorie der modernen Naturwissenschaft* (Springer: Vienna), trans. as *The Logic of Scientific Discovery* (London: Routledge, 2002).

Schervish, Mark J., Teddy Seidenfeld, and Joseph B. Kadane 1984, 'The Extent of Non-Conglomerability of Finitely Additive Probabilities', *Zeitschrift fr Wahrscheinlichkeitstheorie und verwandte Gebiete* 66, pp. 205–26.

—— 2014, 'Non-conglomerability for Countably Additive Measures That Are Not κ-additive', *Review of Symbolic Logic* 10 (2), pp. 284–300.