Fair countable lotteries and reflection

Casper Storm Hansen

Abstract The main conclusion is this conditional: If the principle of reflection is a valid constraint on rational credences, then it is not rational to have a uniform credence distribution on a countable outcome space. The argument is a variation on some arguments that are already in the literature, but with crucial differences. The conditional can be used for either a modus ponens or a modus tollens; some reasons for thinking that the former is most reasonable are given.

According to the canonical axioms for probability formulated by Kolmogoroff (1933), there can be no such thing as a fair countable lottery, that is, a lottery on the set of natural numbers in which each number has the same probability of becoming the winning number. This is because of (an axiom equivalent to) the principle of countable additivity, which states that the sum of the probabilities of a finite or countably infinite family of mutually exclusive events equals the probability of the union over that family. This principle implies that it cannot be the case that the probability of each $n \in \mathbb{N}$ is 0, because that does not add up to 1. In addition, the probabilities of the possible outcomes cannot be the same positive number, because then, the total probability would be larger than 1. And if the probabilities are different, the lottery is not fair.

Therefore, the countable additivity principle implies that there must be a "bias" in favor of small numbers at the expense of large ones. To be precise, it has to be the case for each $\epsilon > 0$ that there is a natural number m such that the probability of the outcome being larger than m is smaller than ϵ. The natural numbers cannot all be treated equally; their position in the natural number sequence has to matter.

For this reason, countable additivity has been objected to by some, notably de Finetti (1972). According to him, the correct interpretation of the probability calculus is a doxastic one. That is, the probability of each event is a measure of the strength of an epistemic agent's belief in the proposition that such an outcome will be realized. Because of his doxastic interpretation of probability, de Finetti believes that the axioms of probability should express the formal constraints on a rational agent's coherent assignment of probabilities, and nothing further. Moreover, he thinks that it is perfectly coherent for such an agent to assign the same probability to each natural number being the winning number in a lottery on $\mathbb{N}$, if she does not have specific information indicating that the lottery is not fair. Say, for instance, that the agent has just been informed that a lottery on the natural numbers is being performed and is given *no* information about how. In that case, de Finetti thinks that it is rationally permissible to distribute the probability uniformly. And that means assigning 0 to each possible outcome (because de Finetti does accept the principle of finite additivity, which is sufficient to rule out a uniform distribution with positive probability for each n).

However, he does not think that distributing the probability uniformly is obligatory. Rejecting the principle of countable additivity is associated with a *permissive* stance regarding rational agents' assignments of credences to a system of propositions. (The extreme other end of the spectrum is the restrictive position that, given any collection of empirical evidence and any set of propositions, there is one, and only one, assignment of credences to those propositions that is consistent with being completely rational.) According to the permissive stance, you should be "allowed" to assign probability 0 to each natural number in the described situation, just as you should be allowed to distribute your credences in a manner that is in line with countable additivity.

The main conclusion of this paper is in the form of a conditional: If the principle of reflection is a valid constraint on rational credences, then it is not rational to have a uniform credence distribution on a countable outcome space. The main argument will be given in section 1. This argument concerns a specific scenario, but section 2 argues that a general lesson can nevertheless be learned from it. The antecedent of the conditional will be made precise in section 4. And the strength of the argument for the conditional is highlighted through comparisons with some similar arguments: one in section 3 and two in section 5.

I have been surprised by the reactions to the argument in section 1. Several people have been convinced that it was invalid. I have therefore included an appendix containing a formal proof of its validity.

The secondary aim of this paper is to map out the dialectical situation to which the conditional gives rise: i.e., that it can be used either for a modus ponens or a modus tollens, depending on the strength of the independent reasons for believing in the antecedent and the negation of the consequent. Reasons for the former are discussed in section 4, and for the latter, in section 6. As will become apparent, I think that a modus ponens, leading to a restrictive stance, is the more reasonable move. However, *this* conclusion is not one I can back up with something as strong as a proof.

1 Main argument

For the time being, assume van Fraassen's (1984) original reflection principle: if an agent at an instant of time t_0 knows that she will, at a later instant of time t_1, have credence x in a given proposition p, then she ought to have credence x regarding p at t_0.

As a warm-up, consider the following banal scenario. Two fair dice, A and B, are rolled. Let N_A and N_B be the stochastic variables for the number of pips. Before seeing the outcome of either, your credence function P_0 satisfies $P_0(N_\mathrm{B} > n) = \frac{6-n}{6}$ for each of the six possible values of n. You then learn that die A shows some specific number of pips, n_A, but remain ignorant about die B's outcome. Your new knowledge that $N_\mathrm{A} = n_\mathrm{A}$ affects your credence in the proposition that die B shows more pips than die A, because you now know that $N_\mathrm{B} > N_\mathrm{A}$ iff $N_\mathrm{B} > n_\mathrm{A}$; you therefore assign the two propositions the same probability. Because of your remaining ignorance and the fact that the dice are stochastically independent,[1] your probability for the latter has not changed. Hence, if P_1 is your credence function at this point in time, then $P_1(N_\mathrm{B} > N_\mathrm{A}) = P_1(N_\mathrm{B} > n_\mathrm{A}) = P_0(N_\mathrm{B} > n_\mathrm{A}) = \frac{6-n_\mathrm{A}}{6}$. The reflection principle does not apply in this case, because this value depends on die A's outcome. Hence, the principle does not imply that you ought to have $P_0(N_\mathrm{B} > N_\mathrm{A}) = \frac{6-n_\mathrm{A}}{6}$. However, if the result had been the same for all outcomes of die A, the reflection principle would have applied. That is what happens when the two dice are replaced with countable lotteries, and you are replaced with an agent who rejects countable additivity.

So now consider the following, interesting scenario. Two stochastically independent lotteries on the natural numbers are conducted. An agent who rejects countable additivity and only accepts finite additivity assigns, for each of the lotteries, probability 0 to each possible winning number, prior to being informed of the actual winning numbers. Let t_0 be an instant of time before the agent is informed of either winning number, but knows what is described in this paragraph. At t_1 the agent is informed of the winning number, n_1, of one of the lotteries. Then at t_2 she is informed of the winning number, n_2, of the second lottery, but simultaneously the knowledge of the first winning number is erased from her memory. She does not have any certainties at t_0 that she doesn't also have at t_1 and t_2 (the relevance of this last stipulation will only become clear in section 4).

Assume that the agent is ideally rational, and accepts both the uniform

[1] That is, for all X and Y such that $N_\mathrm{A} \in X$ and $N_\mathrm{B} \in Y$ are defined, $P(N_\mathrm{B} \in Y) = P(N_\mathrm{B} \in Y \mid N_\mathrm{A} \in X)$, and similarly for A and B reversed. Since zero-probability propositions will be on the table, conditional probability must be understood as primitive (as opposed to defined as a fraction, cf. Hájek 2013), and dictate how an ideally rational agent must update her credences when receiving new information.

distributions for the lotteries and the reflection principle. Let P_0, P_1, and P_2 be her credence functions at t_0, t_1, and t_2, respectively. Additionally, let N_1 and N_2 be the stochastic variables for the two lotteries. At any time t_i when the outcome of the second lottery is not known to the agent, it is the case for all $n \in \mathbb{N}$ that $P_i(N_2 > n) = 1 - \sum_{j=0}^{n} 0 = 1$. This holds in particular for $i = 1$ and $n = n_1$, so $P_1(N_2 > n_1) = 1$. In addition, when the outcome of the first lottery *is* known to the agent, she knows that $N_2 > n_1$ iff $N_2 > N_1$, so she assigns the two propositions the same probability. This also holds in particular for $i = 1$, so $P_1(N_2 > N_1) = P_1(N_2 > n_1) = 1$. By similar reasoning we get $P_2(N_1 > N_2) = 1$, which implies $P_2(N_2 > N_1) = 0$.

As we have been able to deduce this without knowing the winning numbers, the agent is also capable of doing so at t_0. So, from $P_1(N_2 > N_1) = 1$ and the reflection principle, we can deduce $P_0(N_2 > N_1) = 1$. Similarly, $P_2(N_2 > N_1) = 0$ implies $P_0(N_2 > N_1) = 0$. Contradiction.

2 From special case to general principle

So far, it has been established that, in one specific type of scenario, the reflection principle implies the irrationality of having uniform credence functions for certain stochastic variables with countable outcome spaces. It does not follow by logic alone that this holds in general. However, I will argue that the special case nevertheless makes the general principle plausible.

In the field of mathematics, the standard of evidence for a general claim is that a *proof* of that claim is given. That is, however, not the standard in most other fields, including many that make heavy use of mathematics. For instance, we are in many cases happy to accept a general claim in physics merely on the basis of knowing it was correct in the specific case of some observed instances in a laboratory. *Pace* Hume, I believe that this is rational.

Hence, the fact that the field of formal epistemology makes heavy use of mathematics is not sufficient to make demands for *proofs* of general principles reasonable. And we often don't make such demands. A pertinent example is the principle of finite additivity, which many, including de Finetti (1972, section 5.9), accept on the basis of Dutch book arguments. Such arguments are, in the first instance, about a much more specific case than the general principle: namely, those cases in which an epistemic agent assigns credences to a finite number of mutually exclusive propositions *and* in which there is a Dutch bookie present to take advantage of any violations of finite additivity. One *could* claim that these arguments merely demonstrate that it is rationally obligatory to abide by finite additivity in cases where there *actually* is a guarantee of financial loss given failure to abide. Yet, we do not claim that. We implicitly accept that the difference between cases with a bookie and cases without is not substantial with respect to finite additivity; rather, the presence of a bookie merely *reveals* what is the case generally.

Similarly, I believe that the special case considered in the previous section reveals the general irrationality of uniform credence functions on countable outcome spaces. If you are tempted to violate finite additivity, you should consider how you would react if a Dutch bookie were to materialize, and apply the counter-factual conclusion to the actual situation; and if you are tempted to violate countable additivity, you should consider how you would react if a hypnotist with amnestic powers were to materialize and tell you that there is a second stochastic variable similar to and independent of the one you were considering, and apply the counter-factual conclusion to the actual situation.

Consider the ad hoc positions that one is forced into, if one insists that countable additivity is only mandatory in the specific kind of epistemic situation for which a contradiction can be deduced otherwise. Modify the scenario as follows. The agent learns about the first lottery (but not its outcome) at t_{-2}, and adopts a uniform distribution for it. At t_{-1}, she learns about the second lottery, and at t_0, learns the rest of what she was previously stipulated to learn at t_0. Must the agent adopt a non-uniform distribution for the second lottery at t_{-1}? For the agent to *decide* to do so would be unmotivated; for rationality to *demand* it raises the question: why not also the first one? Or is the agent allowed to adopt a second uniform distribution at t_{-1}, but then forced to change at least one of the distributions upon learning the new information at t_0, even though that new information is only about how the outcomes of the lotteries will be learned and forgotten, and does not pertain to the probabilities of the possible outcomes? An opponent would be forced to answer one of these questions affirmatively. That seems just as bad as assigning credence .5 to a proposition A and credence .6 to $\neg A$, and only being willing to revise those assignments to accord with finite additivity if and when it is subsequently revealed that you are about to be Dutch booked, in a case in which that revelation is not relevant to the probability of A.

The reasonable position is that the difference between the agent having one uniform distribution at t_{-2} and having two at t_0 is not substantial, but merely epistemic: irrationality can be *proved* (assuming reflection and finite additivity) for the latter, but not the former. I assume the principle of finite additivity.[2] Therefore, the conclusion is that the reflection principle implies the irrationality of employing a uniform credence function for a stochastic variable with a countable outcome space.

[2]Some deny this principle because they reject the idea that credences can (in ideal cases) be represented by exact probabilities. See, e.g., Fine (1988, 400).

3 Comparison

It is useful to compare the argument in section 1 with another that has been made against the rationality of uniform probability distributions on countable outcome spaces and is also premised on reflection.[3] It concerns a different scenario. First, a fair coin is tossed to determine which of two procedures should be followed to pick a random natural number. If the coin lands heads, the natural number is determined using a method that makes the probability of N being realized as n equal to 2^{-n} for each $n \in \mathbb{N}$. If, instead, the coin lands tails, a method such that the probability is 0 for each n is used. In both cases, an agent is subsequently informed of which natural number was chosen, but not about how it was chosen, i.e., he is not told about the outcome of the coin toss.

The argument concerned with this scenario goes as follows. On the one hand, the agent must assign probability $\frac{1}{2}$ to Heads and to Tails prior to the procedure being carried out. On the other hand, afterward, no matter which outcome n results, he must assign probability 1 to Heads because of Bayes' Theorem, and because the outcome being n has positive probability on the assumption of Heads, but zero probability on the assumption of Tails. The independence of n implies that this can be realized in advance, so by reflection, his credence in Heads before the procedure must equal his credence after. Again we have a contradiction, which—or so the argument goes—can only be avoided by giving up the assumption that it is rationally permissible to employ the uniform distribution in the first place.

In the detailed calculation of the prior probability of Heads using reflection shown below, P and Q are the agent's credence functions before and after, respectively; H is the Heads event; T is the Tails event; N is the stochastic variable for the natural number; and n is the actual outcome for N.

[3] The argument is taken from Howson (2014, section 8). It is only a slight variation on an argument from de Finetti (1972, 205–206) (where it is attributed to Lester Dubins), but the slight variation is exactly that it uses reflection (explicitly), which makes it better suited for comparison with the above argument. The original argument by de Finetti is instead concerned with *conglomerability*, the principle that a probability $P(A)$ is, for any countable partition $\{B_i\}_{i \in I}$ of the outcome space, in the interval spanned by the conditional probabilities $\{P(A \mid B_i)\}_{i \in I}$.

$$P(H) \overset{\text{Reflection}}{=} Q(H) \overset{\text{Conditionalization}}{=} P(H \mid N = n) \overset{\text{Bayes' Theorem}}{=} \frac{P(N = n \mid H) \cdot P(H)}{P(N = n)} \overset{\text{Finite disintegrability}}{=}$$

$$\frac{P(N = n \mid H) \cdot P(H)}{P(N = n \mid H) \cdot P(H) + P(N = n \mid T) \cdot P(T)} = \frac{2^{-n} \cdot \frac{1}{2}}{2^{-n} \cdot \frac{1}{2} + 0 \cdot \frac{1}{2}} = 1$$

Howson provides (again, giving due credit to de Finetti) a defense against this argument. Based on permissiveness, it involves rejecting the equality between $Q(H)$ and $P(H \mid N = n)$. He claims that rational agents have more leeway in deciding how to update their credence functions than is normally assumed by Bayesians. For example, it is supposed to be rationally permissible to pick some arbitrary (but "large") number k and update by letting $Q(H)$ be equal to $P(H \mid N \leq k)$ (which is again equal to 1) if $n \leq k$, and equal to $P(H \mid N > k)$ (which is again equal to $\frac{1}{1+2^k}$) if $n > k$. That way, the value of $Q(H)$ is not a given prior to the procedure, so reflection does not apply.

Insofar as this suggestion has any plausibility, it is due to the fact that when the agent is informed that the outcome is n, that outcome might, as far as he knows, be a result of the sub-procedure for determining the number associated with Tails. In that sense, it might at least seem that this piece of data is epistemically relevant to the agent's assessment of the probability that this sub-procedure had the outcome n.[4] Therefore, it might seem rationally permissible for the agent to assign positive probability to the possibility that the coin came up tails, and that the second sub-procedure delivered the outcome n.

With that in mind, let us return to the new scenario. At t_1 the agent is informed of the outcome, n_1, of the first lottery. Let us say that it is 7. This is definitely information only about the first outcome, and not about the second outcome, as, by assumption, the two lotteries are stochastically independent. So the agent's rational credence for the proposition that the second lottery will have an outcome larger than 7 is not affected by the new information at t_1: we have both $P_0(N_2 > 7) = 1$ and $P_1(N_2 > 7) = 1$. Furthermore, since the agent knows that N_1 is realized as 7, the latter

[4]This will perhaps be clearer if we change the scenario slightly. Assume that no matter the outcome of the coin toss, both of the two sub-procedures are carried out: the stochastic variable N_h has the 2^{-n}-distribution and the stochastic variable N_t has the 0-distribution. And N has the outcome of N_h if heads comes up, while it has the outcome of N_t if tails comes up. Then we can phrase it more clearly: "It might at least seem that this piece of data is epistemically relevant to the agent's assessment of the probability that n was the realization of N_t."

equation implies $P_1(N_2 > N_1) = 1$. To claim that it is rationally permitted for the agent to assign another value to $P_1(N_2 > N_1)$ would be to claim that it is rationally permitted for the agent to cease to consider the probability for N_2 to be distributed uniformly on $\mathbb{N}$, on the basis of new evidence that is irrelevant to N_2.

Of course, the argument to the conclusion $P_1(N_2 > N_1) = 1$ also goes through for any other value of n_1 than 7. Therefore, in this case, reflection does apply, so that $P_0(N_2 > N_1) = 1$. Similarly, we get $P_0(N_2 > N_1) = 0$. Assuming the reflection principle and the rationality of assigning credences uniformly to countable outcome spaces, a contradiction ensues. Even if one accepts de Finetti's and Howson's way out of that conclusion for the scenario they discuss, it cannot be applied here.[5]

A reviewer has pointed out that there is an avenue open for disagreement with this conclusion of mine. They rely on the premise (which I am happy to grant) that for any fixed amount of time between t_0 and t_1, there is an upper bound, k, on the number that can be communicated to the agent as the outcome of the first lottery. Hence, the argument goes, if the outcome exceeds k, the agent is only able to update on $N_1 > k$ and not on $N_1 = n_1$. Then, something close to Howson's instructions for how to update would be justified, without the need for permissive arbitrariness: the agent could be updating on her total new evidence.

I believe that the reviewer's objection succeeds if t_1 and t_2 are fixed instants of time. However, the solution is simply to not consider them as such. Instead, let t_1 be whenever the first winning number has been communicated to the agent, however long that takes, and similarly for t_2. Kierland, Monton, and Ruhmkorff (2008) have investigated whether it matters for reflection whether the instant of time that is reflected on is known to the agent at the time of reflection, and reached the conclusion needed for this solution: it does not.[6]

[5] A popular idea for how to reform probability theory is to allow events to have infinitesimal non-zero probabilities: see Benci, Horsten, and Wenmackers (2018) and references therein. Countable fair lotteries are an important part of the motivation for this move, as infinitesimals allow for the reconciliation of the countable additivity principle with uniformity. However, it does little to address the problem raised here. Since each event of the form $\{0, \ldots, n\}$ in a fair countable lottery is assigned an infinitesimal probability, the above argument goes through if we replace the exact-value version of the reflection principle with this interval version: if an agent at an instant of time t_0 knows that she will, at a later instant of time t_1, assign credence belonging to the interval I to a given proposition p, then she ought to have a credence that belongs to I regarding p at t_0 (and similarly for the other reflection principles to be discussed in section 4).

[6] That instant of time must be *specified* in such a way that the future self will know when it has arrived (Schervish, Seidenfeld, and Kadane 2004). Otherwise, the current self may be able to infer extra information from the condition about the future probability assignment in such a way that it is in a better position to assign probabilities than the future self, and should therefore not defer to it. However, t_1 and t_2 clearly are so specified.

4 Reflection

There is a problem with both the argument in section 1 and the argument in section 3 that has to be rectified. Namely, they are based on a false premise: van Fraassen's reflection principle is not a valid principle of ideal rationality. The literature contains several counter-examples, of which I will give just one here, which is probably the simplest: Talbott's (1991) scenario involving a person who is planning on getting drunk. At the beginning of the evening, when she is still sober, she knows that she will later be drunk and not capable of driving safely. She also knows that when she gets drunk she will be convinced that she can drive safely. It would be irrational for her to reflect on that knowledge of her future credence and adopt the belief, at the beginning of the evening, that it will be safe for her to drive when she is drunk.

Other counter-examples to van Fraassen's reflection principle involving anticipated irrationality can be found in Talbott (1991) and Christensen (1991). Other types of counter-examples involve anticipated memory loss (also Talbott), the mere possibility of memory loss (Arntzenius 2003), and—more controversially—self-locating problems (Elga 2000; Arntzenius 2003).

An element common to all the known counter-examples (and, I believe, *all* counter-examples) is that the agent has reason to consider her future credences untrustworthy. And de Finetti and Howson cannot point to anything implying that the agent's credences at t_1 and t_2 cannot be trusted. By *their* lights, they must be considered to be the kind of trustworthy future credences that one must reflect on to be fully rational—except, of course, that that would lead to contradiction.

Titelbaum (2012, 133) has formulated a different reflection principle that is free of the above-mentioned problems and also sufficient for the argument against fair countable lotteries.[7] Titelbaum's principle says, roughly, that if an agent at an instant of time t knows that she at some instant of time t' rationally assigned or will assign credence x to a given proposition p, conditional on a proposition that is equivalent to the conjunction of all those propositions that she is certain of at t but not at t', then she ought to have credence x regarding p at t.[8] Among other things, this princi-

[7] In fact, the alternative reflection principle is only one element of a comprehensive theory that covers the problematic cases mentioned above. (Another significant element is a generalization of conditionalization.) However, only his reflection principle is relevant for present purposes.

[8] The reason that this statement of the principle is rough is that the word "rationally" has to be defined carefully to yield a precise version of the principle. It should be clear enough what it means in cases like drunkenness, but it is very complicated to explain exactly how to interpret it in the case of self-locating problems. Since that kind of problem is not relevant to the scenarios considered in this paper, I will just refer the interested reader to Titelbaum's book.

ple takes account of the possibility that the agent may forget: a possibility that van Fraassen seems to have idealized away (as is typical in standard Bayesianism). She does not have to defer to a future self, if that future self lacks information that her present self has. However, she does have to defer to those of her own future (trustworthy) credences that are conditional on everything she knows now, but will have forgotten at that future instant of time.

I believe that Titelbaum's reflection principle is correct.[9] However, if this principle in any way seems suspicious on account of using conditional probability in a diachronic rule (like conditionalization, which is rejected by Howson and de Finetti), the following weaker principle—which is implied by Titelbaum's—will also suffice for the argument: if an agent at an instant of time t knows that she at some instant of time t' rationally assigned or will assign credence x to a given proposition p, and she knows that at t' she was or will be certain of all those propositions she is certain of at t, then she ought to have credence x regarding p at t.[10]

Because we stipulated that the agent does not have any certainties at t_0 that she doesn't also have at t_1 and t_2,[11] this weak reflection principle is sufficient to validate the inferences from $P_1(N_2 > N_1) = 1$ to $P_0(N_2 > N_1) = 1$ and from $P_2(N_2 > N_1) = 0$ to $P_0(N_2 > N_1) = 0$. That is, even though memory loss plays a role both in the argument in section 1 and in several counter-examples to van Fraassen's reflection principle, my argument goes through under a reflection principle that takes these counter-examples into account.

It seems reasonable to me to accept this weak principle on account of its intuitive plausibility, i.e., without basing it on anything more fundamental, as long as no clear counter-example has been found. And no clear counter-example has been found. Of course, one might claim that *I* have just found such a counter-example. If so, it is hardly a *clear* counter-example, but since I lack a knockdown argument against that claim, I will refrain from concluding categorically that it is irrational to have a uniform credence distribution on a countable outcome space.

[9] Huisman (2015) argues for saving the permissiveness of mere finite additivity by proposing a weakened form of reflection which does not limit a rational agent's current credence for a proposition to what she knows that credence will be updated to in the future (when she does know that), but only to an interval that is determined by what she would update it to in all of a range of scenarios with *counter-factual* limitations on the knowledge she is going to obtain. I do not see any motivation for this proposed weakening of reflection, which amounts to ignoring the actual knowledge, except as an ad-hoc means to avoid countable additivity.

[10] It seems reasonable to suppose that it is something like this amended principle that Howson (2014, 1006) refers to when he writes "So amended, the principle itself seems sound enough: indeed it would, I believe, be virtually self-contradictory to deny it".

[11] Note that being a certainty cannot be identified with having credence 1, because for each $n \in \mathbb{N}$, we have $P_0(N_1 \neq n) = 1$, even though it is not certain that $N_1 \neq n$.

Let me consider the following objection: there are cases where two experts disagree about the credence for a given proposition, so I *cannot* defer to both of them, and hence I *should* not. The case of the two lotteries is just like that; that the "experts" are future versions of the same agent is inessential. Hence, the scenario tells us nothing about countable additivity, but is rather one among many counter-examples to reflection and similar principles of deference.

My answer is that, while there are indeed cases of disagreeing experts, I very much doubt that there are any that are both relevantly similar and uncontroversially rationally permitted. First, the reason the two experts disagree might be that one or both have employed different priors than me, and if my priors are rationally permissible, I might not be rationally required to defer to them. In our scenario, the credences of the three time-slices of the agent, which I shall call agent$_0$, agent$_1$, and agent$_2$, are based on the same priors, and the reflection principle only applies in such cases,[12] so such an expert scenario would not be a counter-example. Second, the experts may disagree in another scenario in which zero-probability events play an essential role too. Then, we would probably be in the same controversial territory that we are currently occupying, and therefore, that scenario would have little dialectical force. Third, the experts' credences for the proposition in question may not be common knowledge, as they are in our case: agent$_1$ knows that agent$_2$ knows that agent$_1$ knows that $P_2(N_2 > N_1) = 0$, for instance. There are no other cases: if priors are shared, and all relevant ones are positive, and the posteriors are common knowledge, then ideal rationality implies that those posteriors are equal. This follows from a result by Aumann (1976).

In our scenario, agent$_1$ assigns probability 1 to $N_2 > N_1$ because she knows $N_1 = n_1$. Agent$_1$ also knows that agent$_2$ assigns probability 0 to $N_2 > N_1$ because agent$_2$ knows what the outcome of the second lottery is. However, agent$_1$ assigns probability 1 to the reason for $P_2(N_2 > N_1) = 0$ being that agent$_2$ knows the winning number for the second lottery *and* that this winning number is in fact larger than the winning number for the first lottery. Agent$_1$'s knowledge that $P_2(N_2 > N_1) = 0$ does not, therefore, give agent$_1$ reason to reevaluate $P_1(N_2 > N_1)$; and vice versa for agent$_2$. This failure of agent$_1$ and agent$_2$ to reach a consensus even though their credences are common knowledge is, in my view, an anomaly that is indicative of the acceptance of merely finite additivity being misplaced, and not something that is similar in relevant respects to a scenario in which everything is in order.

[12] This is part of the rationality prerequisite for reflection mentioned above; see Titelbaum (2012, 134).

5 More comparisons

Arguments similar to the one I made in section 1 have been considered and rejected by Norton and by de Finetti himself. I will explain why mine is better. Norton (2018, subsection 3.2) writes:

> Consider two [fair] lotteries [on $\mathbb{N}$]. For any outcome on the first, there are only finitely many smaller numbered outcomes on the second, but infinitely many larger numbered outcomes. Therefore the outcome of the second has, with overwhelming probability, the greater number. The same inference, starting with the second lottery machine, concludes that, with overwhelming probability, the outcome on the first has the greater number. Both cannot be true. Therefore an infinite lottery machine is impossible.
>
> The fallacy of this argument lies in set theory, prior to consideration of probabilities. Consider all pairs of natural numbers $\langle m, n \rangle$. For any particular value of m, say M, there are infinitely many $n > M$ but only finitely many $n < M$. It does not follow from this that, for all pairs $\langle m, n \rangle$, there are infinitely many pairs with $n > m$ and only finitely many with $n < m$. The inference from "for any particular value of m" to "for all pairs $\langle m, n \rangle$" requires us to form the union of the sets $\{n : n < M\}$. While this set is finite for any particular M, the union for all M is infinite.

It is, of course, correct that the inference to a contradiction cannot be carried through using set theory alone. You need something else that can justify the move from the particular $P(n > M) = 1$ for each value of M to the general $P(n > m) = 1$. This is not made available by the simple description of the scenario provided by Norton. That something else is available in my scenario, in the form of an epistemic event: when the agent learns that m is realized by M, the two propositions $P(n > M) = 1$ and $P(n > m) = 1$ (using Norton's notation) become equivalent for her.

Here is the relevant quote by de Finetti (1972, 98–99):[13]

> The alleged paradox [. . .] can be stated in the following way: let X and Y be two integers chosen at random [. . .] and independently [. . .]. Then [. . .] given any value x of X, the probability that $Y \leq x$ is zero; analogously, given any y the probability that $X \leq y$ is zero. Thus the probabilities of the events $Y \leq X$ and $X \leq Y$ are both zero but this is absurd since the two events are complementary.

[13] The argument was first made in de Finetti (1930).

> However, it is clear (notice the substitution of x and y with X and Y in the last sentence!) that in the argument it has been implicitly assumed that if the event $Y \leq X$ has zero probability conditional on each of the possible and incompatible hypotheses $X = 1, X = 2, X = 3, \ldots X = x, \ldots$, its (unconditional) probability must also be zero. This property certainly holds for a finite number of hypotheses, but in order to extend it to the infinite case it is necessary and sufficient to assume precisely complete additivity.

While, as de Finetti points out, this argument is also problematic, it is not problematic in the same way as Norton's, which is entirely fallacious and thus establishes nothing. In contrast, de Finetti's argument is, in effect, a sound one for the conditional that has conglomerability as its antecedent and countable additivity as its consequent. However, de Finetti denies both the antecedent and the consequent, and the conditional is dialectically impotent, because it is hard to see why anyone should believe the antecedent if they do not already believe the consequent. Similarly, I presume that de Finetti would have accepted the conditional that I have established, where reflection plays the role of antecedent, while denying its antecedent and consequent. This antecedent, on the other hand, is dialectically interesting. This is because reflection can be tested independently against our intuitions in a wide range of scenarios that are much more down-to-earth and realistic than infinite lotteries, and for which our intuitions are therefore more reliable (see the references in the beginning of section 4). And—at the risk of repeating myself—those tests indicate that the reflection principle can be trusted whenever the credences that are potentially subject to reflection can be trusted.

A possible move is to accept reflection for finite outcome spaces, and deny it for infinite ones. That seems hopelessly ad hoc to me: the strength of the intuition in favor of reflection is not affected by the the size of the outcome space. If I am to be moved, I would need a non-question-begging counterexample to reflection.

6 Permissiveness

As pointed out in the introduction, rejecting countable additivity gives rise to a permissive stance. So does rejecting the weak reflection principle. The latter just happens to be a diachronic principle while the former is a synchronic one; but rejection of either reduces the doxastic obligations on rational agents. The conclusion we can draw from sections 1 and 2 is that the moderate position, i.e., accepting the weak reflection principle while rejecting countable additivity, is unstable: one has to choose between the restrictive stance of accepting both and the extremely permissive stance of rejecting both.

I have given reasons for accepting reflection. But what about de Finetti's most basic intuition for the permissive stance, explained in the introduction of this paper, that if an agent has very little information, it might seem rational for her not to be "biased" in such a way that the position of the possible outcomes in the natural number sequence matters? Well, that intuition leads to inconsistencies by itself, if taken to its full conclusion.

To be completely unbiased concerning the natural numbers, it is not enough to have a credence function P such that for every pair of subsets S_1 and S_2 of $\mathbb{N}$, both of cardinality 1, it is the case that $P(S_1) = P(S_2)$—which is, using a non-standard formulation, what de Finetti wants to allow. This requirement has to be generalized: for every pair of subsets S_1 and S_2 of $\mathbb{N}$ of the same cardinality, it must be the case that $P(S_1) = P(S_2)$. For, if the positions of the elements of S_1 and S_2 in the natural number sequence are disregarded, then there is nothing left to discriminate between the two sets probabilistically, other than the number of elements they contain. And it is easy to see that this requirement is contradictory: using it and finite additivity, we get both

$$P(\{3n+1 \mid n \in \mathbb{N}\}) = P(\{3n+2 \mid n \in \mathbb{N}\}) = P(\{3n+3 \mid n \in \mathbb{N}\}) = 1/3$$

and

$$P(\{3n+1 \mid n \in \mathbb{N}\}) = P(\mathbb{N} \setminus \{3n+1 \mid n \in \mathbb{N}\}) = 1/2.$$

Bartha (2004) attempts to avoid this contradiction by claiming that, in addition to cardinality, the positions and relative distances of the elements of the set have to be taken into account; and that therefore $\{3n+1 \mid n \in \mathbb{N}\}$ and $\mathbb{N} \setminus \{3n+1 \mid n \in \mathbb{N}\}$, because of their different structures, do not have to have the same probability. But that is exactly what de Finetti's intuition opposes: namely, that the positions of the natural numbers have to matter.

Bartha defends his position using an analogy argument. In the case of a uniform distribution on $[0,1]$, the probabilities of any pair of finite subsets is the same (namely 0), while the probabilities of a pair of subsets of cardinality $2^{\aleph_0}$ may be different (e.g., $P([0,1]) = 1$ and $P([0,\frac{1}{2}]) = \frac{1}{2}$). He therefore thinks that we should accept something similar in the case of the natural numbers. I will grant that this argument has some force, but—as is the case with most analogy arguments—it is limited. The outcome spaces $\mathbb{N}$ and $[0,1]$ are very different. The idea that some sort of idealized dart-throwing at a one-meter-long target should be able to result in a uniform distribution on that target has no analogue in the case of $\mathbb{N}$, with its lack of an upper bound on the distances between elements. That two sets with the same cardinality can have different probabilities in a distribution that is considered uniform in the case of an outcome space that is a bounded continuum does not imply, in a straightforward way, that the same should hold in the case of an unbounded, discrete outcome space.

My point is this. *Prima facie* it may seem that de Finetti's principle that a sufficiently ignorant agent should be allowed, without being expelled from the good society of rational people, to be "unbiased" about the elements of $\mathbb{N}$, is one of those basic and almost-obvious truths that one should stick to, almost come-what-may. And if one feels that way, one might be inclined to stick to one's principle, when it is revealed that it conflicts with reflection. But what we have just seen is that the positions and relative distances of the elements of $\mathbb{N}$ have to play a role in *some* contexts—Bartha cannot deny that.[14] It is not possible to treat all possible outcomes completely on a par. So the simple intuition must *at best* be replaced with some more muddy and complicated principle, and muddy and complicated principles are bad candidates for fundamental axioms that can plausibly be claimed to require no further justification. Moreover, if it is acknowledged that further arguments are needed, then are there any that are better than analogy arguments?

The case for (mere) finite additivity seems weak, while the case for reflection seems much stronger. As mentioned earlier, the latter can be—and has been—tested against scenarios in which our intuitions are much more reliable. So, when pitted against each other in a fight, as this paper argues that they are, finite additivity appears to be on the losing side. Or, at least, that is what I would gamble on, if I must.

Appendix

This Appendix contains a longer and more precise version of the argument in section 1.

The first thing to make more precise is the character of the functions. P_0 is stipulated to be a probability function on $\mathbb{N}^2$. Making the dependence of P_1 and P_2 on the outcomes of the lotteries explicit, these are not simply probability functions on $\mathbb{N}^2$, but rather functions from $\mathbb{N}^2$ into the space of probability functions on $\mathbb{N}^2$. For $n, m \in \mathbb{N}$, $s \subseteq \mathbb{N}^2$, and $i = 1, 2$, $P_i(n, m)(s)$ represents the agent's credence at t_i for the proposition that the ordered pair of outcomes belongs to s, if the ordered pair of outcomes is actually (n, m).

In continuity with the notation used above, I will abbreviate "$\{(N_1, N_2) \mid \phi\}$"

[14] The question of how far uniformity can be taken before this problem kicks in is explored by Kerkvliet and Meester (2016) and (using infinitesmal probabilities) Wenmackers and Horsten (2013).

as just "ϕ". So, for instance, $N_2 < l$ is the set $\{(N_1, N_2) \mid N_2 < l\}$, i.e. $\mathbb{N} \times \{0, \ldots, l-1\}$.

Premises

In the statements of the premises, all free variables are implicitly bound by initial universal quantifiers, restricted to

- $\mathbb{N}$ for the variables "n", "m", "k", and "l",
- subsets of $\mathbb{N}^2$ for "s" and "t",
- $[0, 1]$ for "x",
- $\{1, 2\}$ for "i", and
- the set of finite sets of mutually disjoint subsets of $\mathbb{N}^2$ for "S".

The first four premises are uncontroversial principles of probability, including finite additivity:

1. $P_i(n, m)(\bigcup S) = \sum_{s \in S} P_i(n, m)(s)$
2. $P_i(n, m)(\mathbb{N}^2 \setminus s) = 1 - P_i(n, m)(s)$
3. $P_i(n, m)(s) = 1 \rightarrow P_i(n, m)(t) = P_i(n, m)(s \cap t)$
4. $P_0(\mathbb{N}^2) = 1$

The next group of premises encode the scenario stipulations concerning the events at t_1 and t_2:

5. $P_1(n, m)(N_1 = n) = 1$
6. $P_1(n, m)(N_2 = l) = 0$
7. $P_2(n, m)(N_1 = k) = 0$
8. $P_2(n, m)(N_2 = m) = 1$

The final premise comprises the relevant instances of the reflection principle:

9. $P_0(\{(n, m) \mid P_i(n, m)(s) = x\}) = 1 \rightarrow P_0(s) = x$

In fact, this premise is a little more than the reflection principle. The premise does not say that if the t_0 credence of the t_i credence of s being x is 1, then the t_0 credence of s is x; but rather, that if the t_0 credence of an event that leads to the t_i credence of s being x is 1, then the t_0 credence of s is x. This is a strengthening of reflection that is justified when the agent knows which events will lead to which credences. Thus, this formal premise combines what is informally more easily thought of as several premises: the weak reflection principle, the ideal rationality of the agent, and that the agent knows at t_0 what will happen at t_1 and t_2 (except for the specific winning numbers).

Deduction

From 1 and 6:

10. $\forall n, m \in \mathbb{N} : P_1(n, m)(N_2 \leq n) = \sum_{i=0}^{n} P_1(n, m)(N_2 = i) = 0$

From 2 and 10:

11. $\forall n, m \in \mathbb{N} : P_1(n, m)(N_2 > n) = 1 - P_1(n, m)(N_2 \leq n) = 1$

From 3:

12. $\forall n, m \in \mathbb{N} : P_1(n, m)(N_1 = n) = 1$
$$\rightarrow P_1(n, m)(N_2 > N_1) = P_1(n, m)(N_2 > N_1 \cap N_1 = n)$$

Because for all $n \in \mathbb{N}$, $N_2 > N_1 \cap N_1 = n$ is the same set as $N_2 > n \cap N_1 = n$:

13. $\forall n, m \in \mathbb{N} : P_1(n, m)(N_2 > N_1 \cap N_1 = n)$
$$= P_1(n, m)(N_2 > n \cap N_1 = n)$$

From 3:

14. $\forall n, m \in \mathbb{N} : P_1(n, m)(N_1 = n) = 1$
$$\rightarrow P_1(n, m)(N_2 > n \cap N_1 = n) = P_1(n, m)(N_2 > n)$$

From 12, 13, and 14:

15. $\forall n, m \in \mathbb{N} : P_1(n, m)(N_1 = n) = 1$
$$\rightarrow P_1(n, m)(N_2 > N_1) = P_1(n, m)(N_2 > n)$$

From 5, 11, and 15:

16. $\forall n, m \in \mathbb{N} : P_1(n,m)(N_2 > N_1) = 1$

From 16:

17. $\{(n,m) \,|\, P_1(n,m)(N_2 > N_1) = 1\} = \mathbb{N}^2$

From 4 and 17:

18. $P_0(\{(n,m) \,|\, P_1(n,m)(N_2 > N_1) = 1\}) = 1$

From 9 and 18:

19. $P_0(N_2 > N_1) = 1$

By analogous reasoning, using 7 and 8 instead of 5 and 6:

20. $P_0(N_2 > N_1) = 0$

From 19 and 20:

21. Contradiction

Acknowledgements

I have benefitted from discussion with and feedback from Claire Benn, Sharon Berry, Elvira Di Bona, Silvia Jonas, Leora Katz, Timber Kerkvliet, Øystein Linnebo, Olla Solomyak, Georgie Statham, and the anonymous referees.

Conflict of interest statement

The author states that there is no conflict of interest.

References

Arntzenius, F. (2003). Some problems for conditionalization and reflection. *Journal of Philosophy 100*, 356–370.

Aumann, R. (1976). Agreeing to disagree. *The Annals of Statistics 4*, 1236–1239.

Bartha, P. (2004). Countable additivity and the de Finetti lottery. *British Journal for the Philosophy of Science 55*, 301–321.

Benci, V., L. Horsten, and S. Wenmackers (2018). Infinitesimal probabilities. *British Journal for the Philosophy of Science 69*, 509–552.

Christensen, D. (1991). Clever bookies and coherent beliefs. *Philosophical Review 100*, 229–247.

de Finetti, B. (1930). Sulla proprietà conglomerativa delle probabilità subordinate. *Rendiconti del Reale Instituto Lombardo di Scienze e Lettere 63*, 414–418.

de Finetti, B. (1972). *Probability, Induction and Statistics*. Wiley.

Elga, A. (2000). Self-locating belief and the Sleeping Beauty problem. *Analysis 60*, 143–147.

Fine, T. L. (1988). Lower probability models for uncertainty and nondeterministic processes. *Journal of Statistical Planning and Inference 20*, 389–411.

Hájek, A. (2013). What conditional probability could not be. *Synthese 137*, 273–323.

Howson, C. (2014). Finite additivity, another lottery paradox and conditionalisation. *Synthese 191*, 989–1012.

Huisman, L. (2015). Reflecting on finite additivity. *Synthese 192*, 1785–1797.

Kerkvliet, T. and R. Meester (2016). Uniquely determined uniform probability on the natural numbers. *Journal of Theoretical Probability 29*, 797–825.

Kierland, B., B. Monton, and S. Ruhmkorff (2008). Avoiding certain frustration, reflection, and the cable guy paradox. *Philosophical Studies 138*, 317–333.

Kolmogoroff, A. N. (1933). *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Verlag von Julius Springer.

Norton, J. D. (2018). How to build an infinite lottery machine. *European Journal for Philosophy of Science 8*, 71–95.

Schervish, M. J., T. Seidenfeld, and J. B. Kadane (2004). Stopping to reflect. *Journal of Philosophy 101*, 315–322.

Talbott, W. J. (1991). Two principles of Bayesian epistemology. *Philosophical Studies 62*, 135–150.

Titelbaum, M. (2012). *Quitting Certainties: A Bayesian Framework Modeling Degrees of Belief*. Oxford University Press.

van Fraassen, B. (1984). Belief and the will. *Journal of Philosophy 81*, 235–256.

Wenmackers, S. and L. Horsten (2013). Fair infinite lotteries. *Synthese 190*, 37–61.