Fair countable lotteries and reflection

Casper Storm Hansen

Draft of September 24, 2017

Abstract: The main conclusion is this conditional: If the principle of reflection is a valid constraint on rational credences, then so is the principle of countable additivity. The argument for it is a slight variation on two arguments that are already in the literature, but with crucial differences that makes it much stronger. The conditional can be used for either a modus ponens or a modus tollens; some reasons for thinking that the former is most reasonable is given. Finally, a version of the main conclusion with "countable additivity" replaced by "perfect additivity" is considered.

According to the canonical axioms for probability by Kolmogoroff (1933) there can be no such thing as a fair countable lottery, such as, for instance, 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 numbers. 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, some object to countable additivity, notably de Finetti (1972). According to him, the correct interpretation of the probability calculus is a doxastic interpretation. That is, the probability of each event is a measure of the strength of an epistemic agent's belief in the proposition that that 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. And 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 that indicates 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 at the 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 so is the principle of countable additivity. The main argument will be given in section 1 and the antecedent will be made precise in section 2. The conditional 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 2, and for the latter, in section 5. They point to the conclusion that a modus ponens, leading to a restrictive stance, is the more reasonable move.

The strength of the argument for the conditional is highlighted through comparisons with two similar arguments; one in section 1 and another in section 3. In section 4, we will consider, without committing to a conclusion, whether the main argument of this paper can be modified into a valid argument for perfect additivity.

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, assign rational credence x to a given proposition p, then she ought to have credence x regarding p at t_0.

Consider the following 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 value of n_1 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 the next section).

Let P_0, P_1, and P_2 be the functions that describe what the agent's credences ought to be, given acceptance of the uniform distributions for the lotteries and the reflection principle, at t_0, t_1, and t_2, respectively.[1] Furthermore, let N_1 and N_2 be the stochastic variables for the two lotteries. When at t_1 it is known to the agent that n_1 is the realization of N_1, but it is not known to her that N_2 is realized as n_2, the probability of the second outcome being larger than the first is, according to the agent, 1 minus a finite sum of 0's. That is, $P_1(N_2 > N_1) = 1$. At t_2, she knows that N_2 is realized by n_2, but not that n_1 realizes N_1. So by similar reasoning, $P_2(N_1 > N_2) = 1$, which implies $P_2(N_2 > N_1) = 0$.

As we have been able to deduce this without knowing the winning numbers, the agent at t_0, if she is sufficiently rational, is also capable of doing so. 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.

I think it is safe to assume that if there are situations in which it is rational to have a uniform credence function for a stochastic variable with a countable outcome space, then there are also situations in which it is rational to have two stochastically independent uniform credence functions for two such stochastic variables, and it is possible to learn about the outcomes and

[1] To accommodate permissive stances, we do not demand that the three functions are defined for all propositions that are relevant to the lotteries. Instead, the domains are limited to propositions for which there is only one rationally permissible credence at the instant of time in question.

forget them again in the manner described. I also assume the principle of finite additivity, which, as far as I am aware, is universally accepted (notice that this principle is sufficient for the above argument—it does not beg the question by implicitly assuming countable additivity). Therefore, the conclusion is that the reflection principle implies that it is irrational to employ such a credence function.

It will be useful to compare this argument with another argument that has been made against the rationality of uniform probability distributions on countable outcome spaces and is also premised on reflection.[2] It concerns a different scenario. First, a fair coin is tossed to determine which of two procedures should be followed in order 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. By reflection, he therefore also has to assign probability 1 to Heads *prior* to the procedure. Again we have a contradiction, that—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.

The detailed calculation of the prior probability of Heads using reflection is as follows, where 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:

[2]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 $\{P(A|B_i)\}_{i \in I}$. (My argument could also have been premised on conglomerability instead of the reflection principle. However, I think that reflection makes for a more vivid argument, as it is more natural to discuss what an agent's credences should be if they actually get new information than what they should be under supposition. And more importantly, I think it is easier to evaluate reflection independently—see below.)

$$P(H) \overset{\text{Reflection}}{=} Q(H) \overset{\text{Conditionalization}}{=} P(H|N=n) \overset{\text{Bayes' Theorem}}{=} \frac{P(N=n|H)\cdot P(H)}{P(N=n)} \overset{\text{Finite disintegrability}}{=}$$
$$\frac{P(N=n|H)\cdot P(H)}{P(N=n|H)\cdot P(H) + P(N=n|T)\cdot P(T)} = \frac{2^{-n}\cdot\frac{1}{2}}{2^{-n}\cdot\frac{1}{2} + 0\cdot\frac{1}{2}} = 1$$

Howson provides (again, giving due credit to de Finetti) a defense against this argument. It is based on permissiveness and it involves rejecting the equality between $Q(H)$ and $P(H|N=n)$. He claims that a rational agent has more leeway in deciding how to update their credence functions than normally assumed by Bayesians. For example, it is supposed to be rationally permissible to pick some arbitrary (but "large") number k and update by setting $Q(H)$ equal to $P(H|N \leq k)$ (which is again equal to 1) if $n \leq k$ and equal to $P(H|N > k)$ (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.[3] 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 n_1, and not about n_2, 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 implies $P_1(N_2 > N_1) = 1$. To claim that it

[3] 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 as 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".

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 Dubins' scenario, it can not be applied here.[4]

2 Reflection

There is a problem with both my argument and Dubins' that has to be rectified, namely that 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. The simplest example is probably 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).

[4] A popular idea for how to reform probability theory is to allow events to have infinitesimal non-zero probabilities (see Benci et al. (2016) and references therein). Countable fair lotteries is an important part of the motivation for this move, as infinitesimals allow for the reconciliation of the countable additivity principle and uniformity. However, it does little to address the problem raised here. If each event of the form $\{1, \ldots, n\}$ in a fair countable lottery is assigned the same infinitesimal probability, ϵ, then the above argument goes through with the minimal change that 0 has to be replaced by ϵ. If finite events have different infinitesimal probabilities, then the 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 rational 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 2).

An element common to all the counter-examples is that the agent has reason to consider her future credences untrustworthy. (That is definitely common to all the known counter-examples, and I believe that it is common to *all* counter-examples.) And de Finetti and Howson cannot point to anything that implies 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 which is free of the mentioned problems and is also sufficient for the argument against fair countable lotteries. 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.[5] Among other things, this principle 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.[6] 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

[5]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.

[6]In a recent paper, 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 which 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.

ought to have credence x regarding p at t.[7]

Since 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, this weak reflection principle is sufficient to validate the inference 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 $P_0(N_2 > N_1) = 0$.

3 Comparison with a simpler argument

A simplified version of my argument has been considered and rejected by Norton and by de Finetti himself. I will explain why the more complex argument is better. Norton (2017, subsection 3.2) writes the following:

> 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 here. That something else is available in the more complex 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.

[7] 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".

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

> 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.
>
> 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 hypothesis $X = 1, X = 2, X = 3, \ldots X = x, \ldots$, 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.

That is, de Finetti claims that the argument begs the question. In my argument, the question-begging assumption is replaced by reflection. Reflection can be tested independently against our intuitions in a long range of scenarios that are much more down-to-earth and realistic than infinite lotteries and for which our intuitions are therefore more reliable. And (to repeat) those tests indicate that the reflection principle can be trusted whenever the credences that are potentially subject to reflection can be trusted.

4 Perfect additivity

In section 1 we saw that the scenario with two similar lotteries made for a stronger argument than the scenario with the two lotteries with different probability distributions. The former is also more interesting than the latter for another reason, namely that it can be modified to be used in an argument for the principle of perfect additivity: the sum of the probabilities of *any* family of mutually exclusive events equals the probability of the union over that family.[8] That is, the argument shows—if it is sound!—that it is also irrational to consider a lottery on a continuum such as $[0, 1]$ fair. Where the first version of the argument defends orthodox probability theory which includes the principle of countable additivity, this version constitute a challenge to it, as it is common to consider uniform distributions on a bounded

[8] This point was made to me independently by Øystein Linnebo and Timber Kerkvliet.

continuum to be in good standing.[9]

The changes that needs to be made to the scenario are that the two lotteries are on the interval $[0,1]$ instead of $\mathbb{N}$ and that the agent accepts countable additivity (but rejects perfect additivity, which leads to the contradiction). There are also two changes to the argument. The first is that instead of the natural ordering of $\mathbb{N}$, a well-ordering of $[0,1]$ is used. The second is that the agent is committed to $P_1(X_2 > X_1) = 1$ and $P_2(X_1 > X_2) = 1$ because of calculations in which a countable, instead of a necessarily a finite, sum of 0's is subtracted from 1.

The conclusion is that if weak reflection is a principle of ideal rationality, then rational credences must adhere to the principle of perfect additivity. However, there are (at least) two ways out of accepting the conclusion of this argument which were not available in the case of the first argument. I will not take a stand concerning whether they are successful but merely mention the options.

First, it is possible to avoid the conclusion by rejecting the axiom of choice which is needed to ensure that there is a well-ordering of $[0,1]$. The axiom of choice is the most controversial of the standard axioms of set theory, which gives this way of avoiding the conclusion some plausibility: rejecting that axiom might be considered less disruptive to orthodox mathematics than giving up on uniform probability distributions on bounded continua.

Second, there is no language which a human can learn and understand that has a name for each element of $[0,1]$, nor can any specific well-ordering of $[0,1]$ be described in such a language. This means that the agent cannot actually learn the value of x_1 and its position according to the well-ordering at t_1. It is unclear whether this undermines the argument. It might be possible to appeal to a possible infinitary mind and claim that because such an *idealized* agent would be able to learn those things, it is a requirement of *ideal* rationality to abide by perfect additivity.

If the argument for the conditional is sound, then we have the choice between using it for a modus ponens and using it for a modus tollens. A modus tollens would amount to rejecting the weak reflection principle. A modus ponens would mean that standard probability theory for continuous outcome spaces does not codify the requirements of ideal rationality (but if so, it might be the case that standard probability theory gives a model of the requirements of ideal rationality which, although not strictly correct, is a good enough approximation for most purposes).

[9]That a challenge to mere finite additivity is also a challenge to mere countable additivity would not have come as a surprise to de Finetti who thought that of the three options of (mere) finite, (mere) countable, and perfect additivity, "the two extreme ones appear by far the most reasonable since it is more natural to assume that a different behaviour should appear, if at all, in passing from the finite case to the infinite than from the denumerable to the non-denumerable" (de Finetti 1972, 93).

5 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 principle, but rejection of either places fewer doxastix obligations on rational agents. The conclusion we can draw from section 1 is that the moderate position of accepting the weak reflection principle while rejecting countable additivity is unstable: one has to choose between the restrictive stance of accepting both or 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 favor of some natural numbers at the expense of others? 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 the two sets probabilistically than the number of elements they contain. And it is easy to see that this requirement is contradictory: using this requirement and finite additivity we get both

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

and

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

Bartha (2004) attempts to avoid the contradiction with the claim that in addition to cardinality, the positions and relative distances of the elements of the set have to be taken into account, and that $\{3n+1|n \in \mathbb{N}\}$ and $\mathbb{N} \setminus \{3n+1|n \in \mathbb{N}\}$, because of their different structures, therefore 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. If it is correct, as speculated in the previous section, that standard probability for continuous outcome spaces does not correctly reflect the requirements of ideal rationality, then this argument is undermined. If not, I will grant that the 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. 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 which is a continuum does not imply, in a straightforward way, that the same should hold in the case of a discrete outcome space. In the former case, we can appeal to the notion of measure to discriminate, when the notion of cardinality does not. It is not clear that there is a relevant analogue for $\mathbb{N}$.)

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 his 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 rôle in *some* contexts—Bartha cannot deny that. It is not possible to treat all possible outcomes completely on a par. That insight ought to shatter the image of de Finetti's principle as fundamental and obviously true. And weakened like that, it would lose to reflection in a fight. Or, at least, that is what I would gamble on, if I must.

References

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

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 (2016). Infinitesimal probabilities. *British Journal for the Philosophy of Science*, published online.

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

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

Elga, A. (2000). Self-locating belief and the Sleeping Beauty problem. *Analysis 60*, 143–47.

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.

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

Norton, J. D. (2017). How to built an infinite lottery machine. *European Journal for Philosophy of Science*, published online.

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.