ORIGINAL ARTICLE

Why Countable Additivity?

Kenny Easwaran

University of Southern California

It is sometimes alleged that arguments that probability functions should be countably additive show too much, and that they motivate uncountable additivity as well. I show this is false by giving two naturally motivated arguments for countable additivity that do not motivate uncountable additivity.

Introduction

A probability function is required to be non-negative, normalized (tautologies are assigned the value 1), and additive. However, there is a controversy as to how much additivity is required. Additivity principles apply to some collection of incompatible events, and say that the probability of the disjunction of these events must be equal to the sum of their probabilities. All theorists accept this principle in cases where the collection is finite. Kolmogorov notably endorsed the claim that this principle should hold in cases where the collection is countably infinite, though not for larger collections.[1]

Countable additivity is often justified by its fruits—much of the orthodox mathematical theory of probability depends on countable additivity for proofs of various limit theorems and laws of large numbers. However, many philosophers, following Bruno de Finetti, have challenged countable additivity. Although it has some nice mathematical consequences, it also has some bad ones. For instance, it rules out countably infinite fair lotteries.[2] They say that additivity principles should be justified by foundational arguments about the nature of probability, and not by weighing the consequences of the assumption. They say that once we think of things this way, we will see that countable additivity is an arbitrary stopping point—they claim that the justifications for countable additivity extend to full additivity for *all* infinite collections.[3] Since few (if any) theorists endorse full additivity, they say that we should reject the arguments for countable additivity as well, and just endorse finitely additive probability theory.

I will show that this is wrong. I give two arguments that probability functions must satisfy countable additivity, which don't generalize to support full additivity. The first is suggested already by Brian Skyrms (1992) (and in fact, one might see the argument on

Correspondence to: E-mail: easwaran@usc.edu

pp. 255–57 of Spielman 1977 as an example). The second extends an argument of Pruss (2013), who uses a similar principle as an argument against infinitesimal probabilities in countable lotteries. I will not evaluate the consequences of finite, countable, or full infinite additivity here—my purpose in this article is to show that countable additivity is not merely an arbitrary stopping point on the way to full additivity.

I will be considering collections of incompatible events whose disjunction is the necessary truth, which I will call "partitions." It is straightforward to see that violations of an infinite additivity principle for a collection of events whose disjunction is not the necessary truth will lead to violations involving a partition, if one just adds the complement of the disjunction as one more event in the collection. Thus, the restriction to partitions is only a device for simplification, and not a significant restriction of the arguments.

Dutch books

In the context of probability as degree of belief, it is common to use Dutch book arguments to establish constraints on probability. These arguments assume that an agent's degrees of belief are determined by the odds at which she is willing to bet. They then show that for an agent whose degrees of belief fail to satisfy some property, there is a set of bets that she is individually willing to take that together guarantee that she will lose money, while there is no such set of bets for agents whose degrees of belief do satisfy the property.

Jon Williamson (1999) gives such a Dutch book argument for countable additivity in his study. Interestingly, though he doesn't remark on this fact, his argument actually extends beyond countable additivity to full additivity. Importantly, his argument assumes that an agent is willing to accept any bet that occurs exactly at the price considered fair by her degrees of belief. I will rehearse his argument here, with some slight changes of notation. Later, I will show that if we only assume that agents accept bets at strictly favorable prices (as suggested by Skyrms), then the argument only supports countable additivity.

There is a partition A, and for each a in A, the agent has degree of belief q_a in a. The attempted Dutch book will be defined by some stake-maker, who chooses an amount Θ_a to stake on each a in A (subject to a condition to be mentioned shortly). The agent will win Θ_a on this bet if a is true, and will pay $q_a\Theta_a$ for the bet. On the betting interpretation of degree of belief, each of these bets occurs at the price the agent considers to be exactly fair. If Θ_a is negative, then the agent is interpreted as selling a bet. When the bets are resolved, exactly one of them will pay out, but all of the sales will have been transacted. Thus, the net gain for the agent is $\Theta_{a_T} - \sum q_a\Theta_a$, where a_T is the unique member of the partition that happens to be true. There is a Dutch book against the agent iff there is a collection of fair bets whose net result is guaranteed to be negative. That is, there is a Dutch book iff there is some appropriate collection of Θ_a such that each one is strictly less than $\sum q_a\Theta_a$, regardless of which element of A is actual.

The technical condition Williamson assumes on the stakes is that $\sum |q_a\Theta_a|$ is finite. That is, the Dutch book can't rely on an infinite amount of money changing hands just in the initial buying and selling of bets. If the agent both spends infinitely much money buying bets and gets infinitely much money selling bets, then there is no well-defined outcome, so the collection of bets can't really be considered a Dutch book. And as Vann McGee (1999) showed in his article, if one allows the agent to spend infinitely much money buying bets at the beginning, one can create a Dutch book against any agent who allows infinitely many possibilities whose probabilities are nonzero, regardless of what additivity constraint the agent satisfies. There may still be worries about imposing this condition, but the fact that this condition is the same in both Williamson's argument (which supports full additivity) and the modified version (which only supports countable additivity) suggests that this condition is at least not relevant for the issue that I intend to focus on.[4]

The idea of the proof is straightforward. If the q_a sum to more than 1, then the stake-maker has the agent spend more than 1 buying bets of stake 1, so that the agent is guaranteed to lose. If the q_a sum to less than 1, then the stake-maker sells bets of stake 1 on each a to the agent for a total price less than 1, and claims 1 regardless of which a is actual. But if the q_a sum to exactly 1, then the total transaction cost is a weighted average of the stakes, and at least one of the bets must have stakes at least as high as this weighted average, so there is no Dutch book. This argument is made mathematically precise in the following paragraphs.

If there is a_0 with $q_{a_0} < 0$, then there is a Dutch book against the agent—just let $\Theta_{a_0} < 0$ and all other $\Theta_a = 0$. If all q_a are non-negative and $\sum q_a < 1$, then there is a Dutch book against the agent—just set every $\Theta_a = -1$, and then we see that the agent's total outcome is $-1 + \sum q_a$, which is negative, regardless of which a is actually true. If $\sum q_a > 1$, then there is also a Dutch book against the agent—in this case there must (by the definition of an infinite sum as the supremum of the finite sums) be some finite subset $A' \subseteq A$ such that $\sum_{A'} q_a > 1$. Set $\Theta_a = 1$ for a in A', and $\Theta_a = 0$ otherwise. (The restriction to a finite A' ensures that $\sum |q_a\Theta_a|$ is finite.) Then the agent's total outcome will be $\Theta_{a_T} - \sum_{A'} q_a$, which is either $1 - \sum_{A'} q_a$ or just $- \sum_{A'} q_a$ (both of which are negative), depending on whether a_T is in A' or not.

On the other hand, if every q_a is non-negative and $\sum q_a = 1$, then there is no Dutch book against the agent. Let $S = \sum q_a\Theta_a$, which was required to be a well-defined finite real number. Then there is a Dutch book against the agent iff every $\Theta_a < S$. But if every $\Theta_a < S$ then $S = \sum q_a\Theta_a < \sum q_a S$ (because every q_a is non-negative, and at least one is positive, since they sum to 1). But $\sum q_a S = S \sum q_a = S$, which contradicts this inequality. Thus, there is no Dutch book against the agent.

Thus, we have shown that there is a Dutch book against the agent consisting entirely of bets whose price is exactly fair iff there is some partition A such that the q_a are not a collection of non-negative numbers that sum to 1. It is straightforward to see that this condition entails the standard axioms of finitely additive probability theory. However, it also entails countable additivity, and even full additivity. Nothing in this argument requires the set A to be countable.

This Dutch book was already considered by Skyrms (1992). However, he notes that this argument only works if we assume that an agent is willing to accept any bet whose price is exactly fair. He modifies this assumption and shows that the result gives a Dutch book argument for countable additivity, but no more. Note that if $q_a = 0$ and Θ_a is negative, then the exactly fair price is 0. However, this bet gives the agent some chance of losing (if a is true), and no chance of winning any positive amount. Thus, Skyrms doesn't assume that an agent will accept a bet at *exactly* the price equal to the stakes times his/her degree of belief, but merely that an agent will accept a bet at any price *strictly more favorable* than this one.

This suggests a slight modification of the description of the betting formalism. We should require that each bet that is bought or sold is at a strictly favorable price for the agent, rather than an exactly fair one. That is, whenever Θ_a is nonzero (so that a bet is being bought or sold), the agent should also gain some positive amount ϵ_a to make the bet favorable. The stake-maker should choose values (Θ_a, ϵ_a), with $\epsilon_a > 0$ if $\Theta_a \neq 0$, and $\epsilon_a = 0$ if $\Theta_a = 0$, where the agent pays $q_a \Theta_a - \epsilon_a$ for the bet, and wins Θ_a if a is true. For any bet with nonzero stakes, the agent will accept the bet as long as the price is strictly better (by ϵ_a) than the price implied by his/her degree of belief ($q_a \Theta_a$). The condition that only a finite amount of money changes hands amounts to the requirement that $\sum |q_a \Theta_a - \epsilon_a|$ be finite. The Dutch book condition amounts to the claim that regardless of which a_T is the true one, we have $\Theta_{a_T} < \sum (q_a \Theta_a - \epsilon_a)$, which must be a well-defined finite real number.

This modified formalism gives rise to a Dutch book against agents whose degrees of belief violate the traditional probability axioms including countable additivity, but not against agents whose degrees of belief satisfy these axioms while failing to satisfy full additivity for arbitrary infinite sets. The basic idea is that when only a countable collection of bets is needed, then the ϵ_a for these bets can be made small enough that they don't ruin the Dutch book constructed by Williamson's argument. But if the only Dutch books with exactly fair prices are uncountable, then the ϵ_a will add up to an infinite bonus payment for the agent at the beginning, blocking the Dutch book. The remainder of this section goes through the mathematical details of the argument.

Several parts of the Dutch book argument go through with only slight modification. If some $q_{a_0} < 0$, then choose $\Theta_{a_0} < 0$ and $\epsilon_{a_0} < q_{a_0} \Theta_{a_0}$, and choose all other Θ_a and ϵ_a equal to 0, and the result is a Dutch book. If $\sum q_a > 1$, then there is some finite subset A' with $\sum_{A'} q_a > 1$. Let $\epsilon = \sum_{A'} q_a - 1$ and $\Theta_a = 1$ for a in A' and $\epsilon_a = \epsilon/n$ for some n larger than the number of elements in A', and let Θ_a and ϵ_a be 0 for a not in A'. This combination of bets still yields a Dutch book. If every q_a is non-negative and $\sum q_a = 1$, then there is still no Dutch book against the agent, because the positive values of ϵ_a just make things more favorable for the agent.

The case where things differ is the one where the q_a are all non-negative, and $\sum q_a < 1$. If A is countable, then we can still construct a Dutch book. Let $\epsilon = 1 - \sum q_a$ (which is positive) and enumerate the elements of A as a_i. Let each $\Theta_{a_i} = -1$ and $\epsilon_{a_i} = \epsilon/3^i$. Then, regardless of which a_i is true, the agent will end up with

$-1 - \sum \left(-q_{a_i} - \epsilon/3^i\right) = -\left(1 - \sum q_a\right) + \epsilon \sum 1/3^i = -\epsilon + \epsilon/2 = -\epsilon/2$, so the result is a Dutch book.

However, if A is uncountable, every q_a is non-negative, and $\sum q_a < 1$, then there is no Dutch book against the agent using this partition. Since the sum of uncountably many positive numbers is infinite, the facts that every q_a is non-negative and $\sum q_a < 1$ imply that only countably many q_a are nonzero. Together with the facts that $\sum |q_a \Theta_a - \epsilon_a|$ is finite, and that ϵ_a is positive whenever Θ_a is nonzero, this means that there are only countably many Θ_a that are nonzero. If the bets amount to a Dutch book, every Θ_a must be less than $S = \sum (q_a \Theta_a - \epsilon_a)$. Since all but countably many Θ_a are equal to 0, S must be positive. Since the ϵ_a for any bets that are made are all positive, this means that at least some q_a must be nonzero (and thus positive). Since each $\Theta_a < S$, this means that $S = \sum (q_a \Theta_a - \epsilon_a) < \sum (q_a S - \epsilon_a) < \sum q_a S = S \sum q_a < S$, which is a contradiction. Thus, there is no Dutch book.

From these results, we can see that any agent whose degrees of belief satisfy countable additivity (whether or not she satisfies uncountable additivity) will have no Dutch book against her, if she only accepts bets at strictly favorable prices. However, any agent whose degrees of belief violate countable additivity will have a Dutch book against her, consisting only of strictly favorable bets. Thus, if Dutch books with bets at strictly favorable prices give the constraints for degrees of belief, then there is a nonarbitrary reason to accept countable additivity, but not uncountable additivity. It is only if we assume that an agent should accept bets that are priced *exactly* at her degree of belief that we get a Dutch book argument for full additivity.

Comparative probability

Dutch book arguments only make sense on the degree of belief interpretation of probability. The new argument I will give doesn't depend very much on the interpretation of probability—the principle seems equally plausible for chance, degree of belief, logical probability, and many other interpretations that have been proposed. I will use a general principle that should hold for any interpretation of probability to show that probability functions must be countably additive. Of course, the mathematical study of functions that are like probability functions, but merely finitely additive, is still an interesting area of theoretical research, but I claim that to the extent that my principle is plausible for any theory that is appropriate to call "probability," these functions are not probability functions.

The argument will depend on one notion that is not itself part of standard numerical probability theory, which is the notion of comparative probability. If E_1 and E_2 are events, and P_1 and P_2 are probability functions, then I write "$(P_1, E_1) \succ (P_2, E_2)$" to mean that E_1 is (strictly) more likely according to P_1 than E_2 is according to P_2. I assume that this notion is connected to numerical probability, so that if $P_1(E_1) > P_2(E_2)$ then $(P_1, E_1) \succ (P_2, E_2)$, but (as I will discuss later), I don't assume that the converse always holds.

The principle I assume is as follows:

The Comparative Principle: If A is a partition for two probability functions P_1 and P_2, then it is not the case that for every member a of A, $(P_2,a) \succ (P_1,a)$.[5]

This principle is adapted from Argument One of Pruss (2013), where it is used to argue that probabilities should be real-valued and not infinitesimal-valued. To justify it, he says, "surely there could not be a lottery with the same tickets as [another] lottery, and yet still with every ticket being much more likely to win." If there could be two such lotteries, then every ticket-holder would prefer the second lottery to the first, even though the lottery will have only one winner. This would surely be absurd—any change to the probabilities in a single-winner lottery that helps some ticket-holders must make some other ticket-holder worse off. Similarly, if we think of the partition as defining a set of scientific hypotheses that we are uncertain about, rather than lottery tickets, then a violation of this principle would mean that an update from P_1 to P_2 would confirm *every* alternative. Thus, it is reasonable that whenever this condition holds, P_1 and P_2 can't both give possible chances for some lottery on a given set of tickets, or degrees of belief a reasonable agent might have in a given set of scientific hypotheses. So the Comparative Principle seems like a reasonable assumption for many different interpretations of probability.

I will show that if we restrict consideration to real-valued probability assignments, then this principle gives an argument for countable additivity, but not full additivity.

I will illustrate the principle with my argument in favor of countable additivity. de Finetti motivated the rejection of countable additivity by considering the "de Finetti lottery"—one natural number will be drawn uniformly at random, so that every number has the same probability. By finite additivity, that probability must be 0 (if it is real-valued—my argument here is an adaptation of Pruss's argument against infinitesimal probabilities in this same context), because otherwise there would be some finite collection of numbers whose probability added to more than 1, which is impossible. Now consider the "St. Petersburg lottery"[6]—one natural number will be drawn at random, but the distribution is not uniform. Instead, the number will be picked by repeated flip of a fair coin, with the number being the number of flips that occur before the coin first comes up heads. The probability that n is drawn is $1/2^n$.

The partition A is the set of natural numbers—in each probability function, the numbers are considered to be pairwise incompatible, and exhaustive of the possibilities. Now, for a given number n, let us ask in which lottery is this number more likely to come up. In the de Finetti lottery, n has probability 0 (though of course it is possible). In the St. Petersburg lottery, n has probability $1/2^n$. Thus, for each number n, it is strictly more likely to come up in the St. Petersburg lottery than in the de Finetti lottery. These two lotteries have exactly the same possibilities, and yet every possibility is strictly more likely on the second lottery than the first. Thus, the Comparative Principle says that these two functions are not both probability functions. (This argument is exactly parallel to

Pruss's argument against each natural number having an infinitesimal probability, since these values would also all be less than $1/2^n$.)

This can be generalized. Let P be a merely finitely-additive measure function and $A = a_1, a_2, \ldots$ be some countable partition of events such that $\sum P(a_i) < 1$. Let $\epsilon = 1 - \sum P(a_i)$ and P' be such that $P'(a_i) = P(a_i) + \epsilon/2^i$.[7] Then P' is a countably additive measure function where every element of A is strictly more likely than on P, which means that they are not both probability functions.[8]

In each case, all participants in the debate agree that the countably additive function is a probability function. (To deny this, the defender of finite additivity would have to make the far more radical claim that something like the St. Petersburg lottery is not just difficult to create in practice, but in fact impossible as a probability function.) The defender of mere finite additivity insists additionally that the merely finitely additive function is a probability function as well, but this is a violation of the Comparative Principle.

Importantly, this argument only supports countable additivity. For an uncountable partition, it is impossible to strictly increase the probability of every event without there being some n such that at least n events end up with probability greater than $1/n$, which is impossible. Thus, for any pair of probability distributions over an uncountable partition, there must be some event where the first assigns at least as great a probability as the second, and there must be some event where the second assigns at least as great a probability as the first. For example, we might have the following two probability distributions: P_1 is the standard uniform distribution on the interval $[0,1]$, while P_2 is the distribution that gives weight $1/2^n$ to the number $1/n$, and 0 to every interval that doesn't contain any of these numbers.

In order for this to be a violation of the Comparative Principle, it would need to be the case that for every x in $[0,1]$, $(P_1, \{x\}) \prec (P_2, \{x\})$. For $x = 1/n$, this is clear. However, for other values of x, we have $P_1(\{x\}) = P_2(\{x\}) = 0$. We have not yet said anything about how $\prec$ compares an event that has the same numerical probability on two different functions. And it's clear that there must be some cases where $P_1(\{x\}) = P_2(\{x\}) = 0$, and yet strict comparative probability points in one direction or the other. For instance, a case on which x is a possible outcome under P_1 and an impossible one under P_2 would be such a case. So this could conceivably amount to a violation of the Comparative Principle.

But in this case there is a further fact that seems to rule this out. For any x that is not equal to $1/n$, there is some interval containing x that has positive probability on P_1 but probability 0 on P_2. Thus, although P_1 doesn't give a higher numerical probability to any element of this partition than P_2, it does give higher density, so in this case it seems that $(P_1, \{x\}) \succ (P_2, \{x\})$. Thus, some $\{x\}$ are strictly more probable on P_1 than on P_2, making up for the ones that are obviously strictly more probable on P_2 than on P_1. So this pair of probability functions doesn't violate the Comparative Principle. I don't want to assume that probability densities always suffice to determine comparative probability (densities depend on the parametrization of the probability space, but comparative probability presumably doesn't). But in this case it seems clear enough that there is no violation.

   In order to unambiguously apply the Comparative Principle, we would need a case where $P_1(a) < P_2(a)$ for all a in A. But as mentioned above, if A is uncountable, then at least P_2 will violate finite additivity. Thus, the only finitely additive, non-negative, normalized functions that are ruled out as probabilities by the Comparative Principle are the ones that violate countable additivity. This principle does nothing to rule out violations of full additivity. So we have a second argument for countable additivity that fails to extend to full additivity. Countable additivity is not an arbitrary stopping point.

Acknowledgments

Notes

  1 For an infinite collection of non-negative numbers, the sum of that collection is said to be the supremum of the sums of the finite subcollections. That is, it is the smallest number such that no finite subcollection adds up to more than that number. Countable sums are sometimes defined slightly differently, as the limit of a particular sequence of finite partial sums, but as long as the numbers involved are all non-negative (as they are for probabilities), the two definitions will always agree—the order only matters to the standard sum definition when negative terms are involved. As it turns out, an infinite collection of non-negative numbers has a finite sum iff there is a countable subcollection such that these ones have a finite sum (equal to the sum of the whole collection), and all other elements of the whole collection are 0. To see this, just note that an uncountable collection of nonzero values has an infinite sum, and values that are 0 make no contribution to the total sum.
  2 A fair lottery should have the same probability for each outcome. For an infinite lottery, even finite additivity means that this probability can't be positive, or else they would sum to more than 1. Thus, if there is a countably infinite fair lottery, then all outcomes would have probability 0, which would violate countable additivity, since the sum of any number of 0s is 0.
  3 This requirement of full additivity amounts to the requirement that the probability distribution be "discrete," in the technical sense that there are at most countably many possibilities with non-zero probability, and these add up to 1. Countable additivity without full additivity additionally allows distributions, traditionally called "continuous," where there are uncountable collections of possibilities that individually have probability 0 but collectively have nonzero probability. Finite additivity without countable additivity allows even more distributions, like de Finetti's countably infinite fair lottery.
  4 I'd like to thank Andrew Bacon for pressing me on this worry.
  5 This principle has some similarities to a principle known as "Conglomerability," though it is in a sense a dual to it. Conglomerability is the claim that for any partition A, there is no event E such that $P(E|a) > P(E)$ for all a in A. If we assume that the conditional distribution given an event is itself a probability function whenever we start with a probability function, then the Comparative Principle entails that there is no event E such that $P(a|E) > P(a)$ for all a in A, which is similar to an instance of Conglomerability, but distinct from it. As it turns out,

Conglomerability itself has some connection to countable additivity. Although it is known that Conglomerability for all countable partitions is equivalent to countable additivity (Schervish, Seidenfeld, and Kadane 1984), it is also known that Conglomerability for uncountable partitions leads to further problems (Kolmogorov 1950, Ch. 5).

6 I give it this name because of its similarity to the definition of the classic "St. Petersburg game" of decision theory, which is not relevant here.

7 We may need to do some more work to extend P' to every event. For events that are unions of events from this partition, we can extend it just by applying countable additivity. If there are events that cross-cut the members of this partition, then some of the a_i must themselves be unions of smaller events, and we have to divide the excess probability among these smaller events. I suspect that the decomposition in (Yosida and Hewitt 1952) and (Schervish, Seidenfeld, and Kadane 1984) of every finitely additive probability function into a purely finitely additive one and a countably additive one will do the work that is required.

8 There will be some events that are more likely on P than on P'. For instance, let E be the complement of one of the a_i. Since $P(a_i) < P'(a_i)$ we have $P(E) > P'(E)$. However, this is of no help—E is an event that is strictly less likely on P', but every outcome in E is strictly more likely on P', which seems like an even more problematic violation of a related comparative principle.

References

Kolmogorov, A. N. *Foundations of the Theory of Probability*. New York: Chelsea, 1950.

McGee, V. "An airtight Dutch book." *Analysis* **59**.4 (1999): 257–65.

Pruss, A. "Infinitesimals Are Too Small for Countably Infinite Fair Lotteries." *Synthese* (2013). doi: 10.1007/s11229-013-0307-z.

Schervish, M. J., T. Seidenfeld, and J. B. Kadane. "The Extent of Non-Conglomerability of Finitely Additive Probabilities." *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* **66** (1984): 205–26.

Skyrms, B. "Coherence, Probability and Induction." *Philosophical Issues* **2** (1992): 215–26.

Spielman, S. "Physical Probability and Bayesian Statistics." *Synthese* **36** (1977): 235–69.

Williamson, J. "Countable Additivity and Subjective Probability." *The British Journal for the Philosophy of Science* **50** (1999): 401–16.

Yosida, K. and E. Hewitt. "Finitely Additive Measures." *Transactions of the American Mathematical Society* **72**.1 (1952): 46–66.