Brit. J. Phil. Sci. **55** (2004), 301–321, axh206

Countable Additivity and the de Finetti Lottery

Paul Bartha

ABSTRACT

De Finetti would claim that we can make sense of a draw in which each positive integer has equal probability of winning. This requires a uniform probability distribution over the natural numbers, violating countable additivity. Countable additivity thus appears not to be a fundamental constraint on subjective probability. It does, however, seem mandated by Dutch Book arguments similar to those that support the other axioms of the probability calculus as compulsory for subjective interpretations. These two lines of reasoning can be reconciled through a slight generalization of the Dutch Book framework. Countable additivity may indeed be abandoned for de Finetti's lottery, but this poses no serious threat to its adoption in most applications of subjective probability.

1 Introduction

The axiom of countable additivity (CA) plays a critical role in modern probability theory. The axiom states:

(CA) If we have a countable infinity of outcomes H_1, H_2, ... that are mutually exclusive, then $P(H_1 \vee H_2 \vee \ldots) = P(H_1) + P(H_2) + \ldots$.

Kevin Kelly ([1996]), following de Finetti ([1974]), questions whether (CA) is an indispensable constraint on subjective interpretations of probability. In such interpretations, particularly as applied to the justification of scientific

© British Society for the Philosophy of Science 2004

hypotheses, (CA) assumes great epistemological significance because of its role in deriving the convergence theorems.[1] In essence, (CA) forces us to adopt the biased view that if there is ever going to be a counterexample to a universal hypothesis, we are far more likely to find it in some finite segment of the future than in the entire remainder of history. For this reason, Kelly ([1996], p. 323) believes that the principle 'should be subject to the highest degree of philosophical scrutiny', rather than being adopted purely for its mathematical merits. He goes on to present instances in which, he alleges, our epistemic intuitions give us good grounds to reject countable additivity.

The opposing point of view, that (CA) is not significantly more problematic for the subjective interpretation than any other axiom of the probability calculus, is represented by people such as Howson and Urbach ([1993]) and Williamson ([1999]). Their main argument is that a Dutch Book justification can be given for (CA), just as for any of the other standard axioms. Dutch Book arguments can be criticized,[2] but I do not wish to call them into question in this paper. One of my main objectives is to show that, even within a 'Dutch Book' framework for subjective probability, we can make sense of Kelly's objections.

A crucial and much-discussed test case for these opposing positions is an example due to de Finetti, which (in slightly altered form) I refer to as the *de Finetti lottery*. Section **2** describes the de Finetti lottery. Section **3** evaluates the Dutch Book argument, and one additional line of reasoning, that countable additivity is mandatory for this example. Section **4** argues, to the contrary, that (CA) may reasonably be rejected for the de Finetti lottery without abandoning the idea that Dutch Book arguments have normative implications for subjective probability—though for this argument to succeed, we have to move to a slightly more general formulation of Dutch Book arguments. Section **5** discusses, and resolves, a paradox that seems to be generated by the ideas in Section **4**.

Section **6** argues that the (plausible) failure of (CA) in the case of the de Finetti lottery does not support broad scepticism about it. It is true that finding a justification for dropping the axiom of countable additivity in a particular case demonstrates that it is not quite so fundamental as the other axioms of the probability calculus. Nevertheless, the exceptional cases can be identified and quarantined. They pose no real threat to the use of (CA) in applications of subjective probability to the philosophy of science.

2 The de Finetti lottery

De Finetti ([1974]) claimed that we should be able to make sense of a uniform probability distribution over a countably infinite set, such as natural numbers.

[1] Savage ([1972]) and Edwards, Lindman and Savage ([1963]) are standard sources.
[2] For example, see Maher ([1997]).

To make this vivid, let us imagine a lottery in which the number of tickets issued is countably infinite, one for each positive integer, and each ticket has an equal (subjective) probability of winning. Such a lottery appears to be conceivable. The assumption that each ticket is equally good seems reasonable, or at least not *a priori* false.

As de Finetti pointed out, however, there is no way to assign an equal probability to each ticket's winning if we accept countable additivity. If we let p_n be the probability assigned to ticket n, then the numbers p_n have to satisfy two conditions:

$$\text{(Equiprobability)} \qquad p_n = p_m \text{ for all } n, m$$

$$\text{(Countable additivity)} \quad p_1 + p_2 + p_3 + \ldots = 1$$

Equiprobability is just the desired assignment of a uniform probability distribution over all tickets. Countable additivity tells us that the probability that *some* ticket wins (which is 1) is the infinite sum of the probabilities that each individual ticket wins. The two conditions cannot, however, both be satisfied. If each $p_n = 0$, the infinite sum will be 0, but if each p_n is the same positive number α, then the series diverges. Countable additivity compels us *a priori*, as de Finetti ([1972], pp. 91–2) says, 'to assign practically the entire probability to some finite set of events, perhaps chosen arbitrarily'. He finds this deeply puzzling:

> What is strange is simply that a formal axiom, instead of being *neutral* with respect to evaluations [...] and only imposing formal conditions of coherence, on the contrary, imposes constraints of the above kind without even bothering about examining the possibility of there being a case against doing so. ([1974], p. 122)

De Finetti abandons countable additivity and retains equiprobability by letting each $p_n = 0$. His argument rests on intuitions about symmetry. There is a partition into a countable infinity of outcomes: exactly one ticket will win. Since there is no further knowledge, we must (or at least we should be permitted to) regard any two tickets as interchangeable, for any two ticket-holders are in an epistemically indistinguishable position.[3] Note that the fundamental intuition here is that the probabilities are *equal*, not that they are all 0. The choice $p_n = 0$ appears to be forced upon us, however, if we want uniformity.

Similar intuitions support a uniform probability assignment in two analogous situations: a finite lottery, and a lottery over the real numbers in the interval [0, 1]. In these analogous lotteries, equiprobability is perfectly

[3] Williamson ([1999]) makes a convincing case that one is not compelled to adopt a uniform prior probability distribution. Like him, I am most concerned with the tenability of the position that one *may* do so. For a very different perspective on what ignorance entails in this type of situation, see Jeffreys ([1961], pp. 119–125).

reasonable and unproblematic (for there is no conflict with countable additivity). Why can't we retain it in the case of a countably infinite set?

3 Two objections to equiprobability

I think that de Finetti is right not to give up on equiprobability, but wrong to let each $p_n = 0$. Sections **3.1** and **3.2** present two ultimately inconclusive arguments for keeping countable additivity and dropping equiprobability. Section **4** shows how we can retain equiprobability without directly rejecting countable additivity—a mysterious-sounding claim, but the mystery will be resolved shortly.

3.1 The 'No random mechanism' argument

The most popular direct objection to a uniform probability assignment in the de Finetti lottery (Spielman [1977]; Howson and Urbach [1993]) is that any mechanism for choosing a positive integer (i.e., a particular ticket number) will inevitably yield a biased distribution. As Howson and Urbach ([1993], p. 81) write: 'it is not at all clear what selecting an integer at random could possibly amount to: any actual process would inevitably be biased toward the "front end" of the sequence of positive integers.'

Even if we concede this point,[4] the natural response (Williamson [1999]) is that it applies only to physical chance, not to subjective probability. There is no need to conceive of a physical mechanism that could select each integer with equal probability. Since we are dealing with subjective probability, we can abstract away from the mechanism used to select the winning ticket. Indeed, the point de Finetti wants to make could be expressed in this way: lacking any precise knowledge about the mechanism, the only reasonable thing to do is to regard each ticket as equally likely to win.

If we are not persuaded by this response, it might be because we believe that subjective probabilities should reflect our knowledge of physical chances. One familiar attempt to represent this connection is Lewis' Principal Principle (Lewis [1980]):

(PP) $\text{Prob}(A/P(A) = r) = r$

Your subjective probability for A, given that the chance of A is r and no other inadmissible information (e.g., that A is or is not true), is r. We might expect, similarly:

(PP′) $\text{Prob}(A/P(A) \neq r) \neq r$

[4] But see McCall and Armstrong ([1989]), which proposes an unbiased mechanism of sorts: God is conducting the lottery—presumably in a fair manner!

Or, with still greater generality, what we might call the Unprincipled Principle:

(UP) Given that the actual physical chance distribution cannot have certain features, our subjective probability distribution ought not to have those features.

If we know that there is no physical mechanism that gives each integer an equal chance to be selected, then our subjective probabilities should not be uniform.

Although there might be some way to salvage this argument, it will not be via UP. Even the special case PP′ is wrong. Choose randomly between two superficially indistinguishable coins, one with a bias $P(\text{heads}) = 0.9$, and the other with a bias $P(\text{tails}) = 0.9$. Our subjective probability for heads should be 0.5, despite knowing that the chosen coin is not fair. We obtain this subjective probability as a weighted average of the two probability distributions we could have if we knew which coin we had chosen.

It might be objected here that given the full set-up (including the random choice), $P(\text{heads})$ is really 0.5. But if this point is conceded, then a parallel move is available for the case of the lottery. A similar intuition about averaging supports equiprobability, although it is not so easy to formulate. For every mechanism M that favours integer m over integer n, there will be another mechanism M' that is identical except that the values p_m and p_n are exchanged. (To construct M', just do a swap and then apply M.) The full set-up of the lottery includes the choice of a mechanism. So our subjective probability assignments should ensure that $p_m = p_n$ for any m and n. The fact that the particular mechanism selected is inevitably non-uniform does not force us into non-uniform subjective probabilities, just as with the biased-coins example.

What we should conclude from all of this is that the arguments are not decisive for either side in the debate. We have merely shifted the dispute to the higher-level problem of weighting distributions. Those inclined to reject equiprobability in the original lottery will reject the preceding paragraph's appeal to symmetry at the level of distributions; those sympathetic to invoking symmetry at the higher level will be sympathetic to it in the original problem.

There is a history of paradoxes that attend such symmetry-based arguments, and it might be argued that this breaks the stalemate in favour of a non-uniform assignment. We know that symmetry arguments can sometimes yield inconsistent results. But sometimes they work. Harmless applications of uniformity require considerable analysis and defence, but they exist. Sections **4** and **5**, together with the Appendix, provide a rigorous justification for the appeal to symmetry in the de Finetti example. The symmetry-based approach to this particular problem remains a viable option.

3.2 The Dutch Book argument

There is a straightforward Dutch Book argument for countable additivity as a constraint on subjective probability which, if successful, appears to rule out a uniform probability assignment in the de Finetti lottery.[5] The argument is a simple generalization of the usual Dutch Book argument for finite additivity.

Suppose that p_i is our fair betting quotient for the proposition that ticket *i* wins, and 1 is our fair betting quotient for the proposition that some ticket wins. Suppose that countable additivity is violated, so that

$$p_1 + p_2 + \ldots < 1.$$[6]

Each of the following bets is fair: bet *against* ticket *i* with a stake of \$1 and betting quotient p_i. This bet pays p_i dollars if ticket *i* loses (we win our bet), and $(p_i - 1)$ dollars if ticket *i* wins (we lose our bet). The system consisting of all these bets taken simultaneously is fair. Suppose now that ticket N wins, as must happen for some N. We win p_i for all tickets other than N, and $p_N - 1$ for ticket N. Our net gain is therefore

$$(p_1 + p_2 + \ldots) - 1$$

which, by assumption, is negative. So no matter what happens, we lose money. This constitutes a Dutch Book.

This argument shows that *if* we assign a standard real-valued betting quotient to the proposition that ticket *i* wins, for each *i, then* these betting quotients must sum to 1 on pain of vulnerability to a Dutch Book. In particular, de Finetti's own solution, which sets $p_i = 0$ for each *i*, is unacceptable.

There are, however, at least two ways to avoid the conclusion that countable additivity is forced upon us. One is to use non-standard probabilities. Assign to each proposition that ticket *i* wins an equal but infinitesimal degree of belief. This may be accomplished in such a way that the *hyperfinite* sum of these probabilities is 1 even though any finite sum is infinitesimal.[7] Arguably, the total net gain should also be computed using hyperfinite summation rather than standard countable summation. But I will not dwell on this point because I want to focus on an alternative approach.

The second way to avoid the conclusion that countable additivity is rationally required is just to deny that we *have* a real-valued degree of belief, or fair betting quotient, for the proposition A_i that ticket *i* wins. Any positive number is too large (given our desire to assign the same betting quotient to each ticket

5 Here I draw on Howson and Urbach ([1993]) and Williamson ([1999]), but such arguments go back to Jeffrey ([1957]). See Skyrms ([1984]) for a review of early examples.

6 This sum can never exceed 1, by finite additivity.

7 For one way of doing this, see Bartha and Hitchcock ([1999]).

number), while 0 is too small (the bet would cost nothing and might pay off). On this view of things, the Dutch Book argument never gets off the ground. We might have a betting quotient of 0.5 for 'an even-numbered ticket wins'. We might also have, as explained in the next section, a *relative* betting quotient of 1 for the pair of propositions A_i (ticket i wins) and A_j (ticket j wins), reflecting the idea that the two propositions are equally likely to be true. The crucial point, though, is that we might lack betting quotients for these propositions taken in isolation. If we have no betting quotients for these propositions, and hence no subjective probabilities, then countable additivity is inapplicable rather than violated. This is not the bland observation that if there are no betting quotients, there are no subjective probabilities at all. Rather, we shall see that it is possible to define a type of betting quotient that is faithful to all of the standard axioms of the probability calculus except countable additivity. The next two sections flesh out and defend the viability of this approach.

4 Equiprobability and relative betting quotients

This section shows that a relationship of equiprobability between two outcomes (or two propositions expressing sets of outcomes) can be defined independently of the existence of any probability function. We define the relationship in terms of *relative betting quotients* and then apply it to the de Finetti lottery. A relative betting quotient for a pair of propositions tells us, roughly speaking, how to trade off against each other a bet for one and a bet against the other.

The (fair) betting quotient for A is a real number p between 0 and 1 such that a bet on A that costs pS and pays S if A is true and nothing if A is false is subjectively fair, for any stake S. To define relative betting quotients, first consider a special case. If two outcomes A and B have well-defined betting quotients p and q respectively, and $p \neq 0$, then the relative betting quotient of B compared to A, written RBQ(B; A), is just q/p. Suppose that this ratio is k, so that (informally) we consider outcome B to be k times as likely as outcome A. Table 1 represents a bet *for A* with stake k and a simultaneous bet *against B* with stake 1.

This system of bets is subjectively fair, and the betting quotients p and q disappear in the final column. Only the ratio, k, matters for the net payoff. No

Table 1 Betting quotients

A	*B*	*For A (stake k)*	*Against B (stake 1)*	*Net gain*
F	F	$-pk$	q	0
T	F	$(1-p)k$	q	k
F	T	$-pk$	$-(1-q)$	-1
T	T	$(1-p)k$	$-(1-q)$	$k-1$

Table 2 Relative betting quotients

A	*B*	*Payoff to the agent*
F	F	0
T	F	kS
F	T	$-S$
T	T	$(k-1)S$

money changes hands if neither A nor B is true. We now generalize this idea to encompass cases where A and B lack betting quotients or (in some cases) where both $p=0$ and $q=0$.

An agent's *relative betting quotient* for B relative to A, written RBQ(B; A), is a non-negative real number k such that the bet described by Table 2 is subjectively fair for any stake S.

If there is no such unique k, then RBQ(B; A) is undefined.

The idea is simple. If neither A nor B is true, no money changes hands. If A is true, the bookie pays out kS. If B is true, the agent pays the bookie S. Note that the stake S may be negative, in which case the direction of gains and losses is reversed. The agent regards a bet on A with payoff kS to be of equal value to a bet on B with payoff S. If A and B have well-defined betting quotients p and q with p non-zero, k is just the ratio q/p.[8]

There are two special cases. If A is a contradiction but B is not, then no value of k makes the bet fair, so that RBQ(B; A) is undefined. If both A and B are contradictions, then any value of k makes the bet fair (since no money will ever change hands in any case), so once again we leave RBQ(B; A) undefined.

We can define a Dutch Book for a system of relative betting quotients in exactly the same way as for ordinary betting quotients. The following results may be derived for any system of relative betting quotients that is not vulnerable to a Dutch Book in a manner parallel to the usual Dutch Book arguments:[9]

(R1) (*Positiveness*) RBQ(B; A) ≥ 0 whenever defined.

(R2) (*Reflexivity*) RBQ(A; A) $= 1$ for any A that is not a contradiction.

(R3) (*Tautologies*) If T is a tautology, then RBQ(A; T) ≤ 1 whenever defined.

(R4) (*Contradictions*) If $\perp$ is a contradiction and A is not, then RBQ($\perp$; A) $= 0$.

(R5) (*Finite additivity*) If B and C are mutually exclusive and RBQ(B; A) and RBQ(C; A) are defined, then RBQ($B \vee C$; A) is defined and RBQ($B \vee C$; A) $=$ RBQ(B; A) $+$ RBQ(C; A).

8 Note that a relative betting quotient may be defined even if both p and q are 0. For instance, if A is the proposition that I am on the moon one second from now, and B is the proposition that I am on the dark side of the moon one second from now, we might reasonably assign RBQ(B; A) $= 0.5$ (or some value close to this). But we might have no defined relative betting quotient for this case.

9 Compare Bartha and Johns ([2001]), where the formalism is slightly different.

(R6) (*Generalized conditionalization*) If RBQ($B; A$) and RBQ($C; A$) are defined and RBQ($B; A$) $\neq 0$, then RBQ($C; B$) is defined and RBQ($C; B$) = RBQ($C; A$)/RBQ($B; A$).

These results are analogues of familiar properties of probability. We do not, however, get countable additivity of relative betting quotients—at least not in general. Unlike simple betting quotients, relative betting quotients have no upper bound. There is no reason why an infinite sum of relative betting quotients should converge at all. Though in the special case where RBQ(B_n; A) is defined for all n, the B_ns are exclusive, and RBQ(B; A) is defined where B is the infinite disjunction of the B_ns, the Dutch Book argument of Section **3.2** can be adapted to show that

$$(\mathbf{CA}^*) \quad \sum_{n=1}^{\infty} RBQ(B_n; A) = RBQ(B; A)$$

If RBQ(B; A) is defined and the assignment of relative betting quotients is coherent, we call this value a *relative probability*, which we write as R(B, A). We are most interested in the case where R(B, A) = 1, in which case we say that A and B are *equiprobable*. If we exempt contradictions, then the relationship of equiprobability is reflexive (by (R2)), symmetric (by (R2) and (R6)) and transitive (again by (R6) and (R2)); hence, it is an equivalence relation.

If R(B, X) is defined, where X represents the entire outcome space, then write $Pr_R(B)$ for this value. It may happen that Pr_R is a (countably additive) probability function. In this case, we say that Pr_R is the *induced* probability measure (i.e., induced by R). $Pr_R(E)$ may not be defined for every proposition E describing a set of outcomes.

Return to the de Finetti lottery. We are finally in a position to assert that any two propositions of the form 'ticket i wins' and 'ticket j wins' are equiprobable, because their *relative* betting quotient (and hence their relative probability) is 1 even though they have no well-defined betting quotients (and hence no induced subjective probability value).

With this analysis in hand, we see that our original question about the necessity of countable additivity should be answered negatively in one sense and affirmatively in another. If the issue is whether all subjective probabilistic reasoning must always be constrained by countable additivity (assuming that Dutch Book arguments are normative for subjective probability), then the answer is negative. On this point, de Finetti and Kelly are correct, though the justification provided here is novel: it might be that we have only relative subjective probabilities, so that countable additivity just does not apply (though analogues of the other axioms still hold). If the issue is whether subjective probabilities properly speaking must satisfy countable additivity whenever they are defined, then (contrary to de Finetti and Kelly) the answer still appears to be affirmative, granting the Dutch Book argument of Section **3.2**. The de Finetti lottery does not provide a counter-example.

To appreciate this point, consider equation (CA*) above, letting B_n stand for 'ticket n wins' and A stand for 'some ticket wins', i.e., the entire outcome space. If each relative betting quotient RBQ(B_n; A) were 0, we would have our original puzzle all over again. Each B_n would have induced probability 0, and we would have a violation of countable additivity for the induced probability function. The alternative approach made possible by the foregoing analysis is that none of these relative betting quotients is defined! In more colourful language: the propositions 'ticket n wins' and 'some ticket wins' are probabilistically incommensurable.

In short, two ideas identified as incompatible in Section **2**—equiprobability over a countable partition and countable additivity (based on Dutch Book arguments)—are in conflict only if we assume that we have well-defined betting quotients (and hence well-defined subjective probabilities). Where we have only relative betting quotients (and relative probabilities), the incompatibility disappears because, although Dutch Book arguments still apply, they do not inevitably yield countable additivity.

5 The re-labelling paradox

5.1 The paradox

Our solution in Section **4** is threatened by the following example, which purports to show that positing a relationship of equiprobability between sets of outcomes in a countably infinite set leads to paradox.[10]

1. Let A be a countably infinite population. Label its members $a_1, a_2, a_3, \ldots$. One individual a_n is to be selected. Suppose, for *reductio*, that any two individuals are equally likely to be selected.

2. Assuming that the relative probability function R is well-defined and induces a probability function Pr_R, we should have $Pr_R(\text{EVEN}) = Pr_R(\text{ODD}) = \frac{1}{2}$
where EVEN $\equiv$ selected a_n has an even label,
and ODD $\equiv$ selected a_n has an odd label.

[10] The example is due to John Norton (in correspondence). Similar sorts of paradoxes depending on a choice of parameter have long been familiar; see Jeffreys ([1961], p. 372). Jaynes ([1985], p. 123) cautions us against 'rushing headlong into a sticky swamp of paradoxes' that arise when approaching infinite sets.

3. Similarly, we should have $Pr_{\mathrm{R}}(\mathrm{ONE}) = Pr_{\mathrm{R}}(\mathrm{TWO}) = Pr_{\mathrm{R}}(\mathrm{THREE}) = Pr_{\mathrm{R}}(\mathrm{FOUR}) = \frac{1}{4}$
where ONE $\equiv$ selected a_{n} has a label of the form $n = 4m + 1$
TWO $\equiv$ selected a_{n} has a label of the form $n = 4m + 2$
THREE $\equiv$ selected a_{n} has a label of the form $n = 4m + 3$
FOUR $\equiv$ selected a_{n} has a label of the form $n = 4m$.

(Proposition ONE is true if the individual selected is one of a_1, a_5, $a_9, \ldots$; proposition TWO is true if the individual selected is one of a_2, a_6, $a_{10}, \ldots$, and so forth.)

It seems reasonable that each of the propositions ONE, TWO, THREE, FOUR, which all express the idea that the selected individual belongs to a set consisting of every fourth member of the original list, should have relative betting quotient (and hence relative probability) of $\frac{1}{4}$ with respect to the full set of outcomes.

4. The way we label the members of the population A should not affect these probabilities. Let us re-label the population, in effect re-arranging the position of the members of the original list to produce a new list $b_1, b_2, \ldots$. First, spread out the even items a_2, a_4, ... so that they now occupy slots 4, 8, 12, ... in the new list. Second, into slots 2, 6, 10, ... of our new list, put items a_3, a_7, $a_{11}, \ldots$, i.e., set THREE. Finally, fill all of the odd slots of our new list with a_1, a_5, $a_9, \ldots$, i.e., set ONE. The resulting list looks like this:
$a_1, a_3, a_5, a_2, a_9, a_7, a_{13}, a_4, a_{17}, a_{11} \ldots$
So $b_1 = a_1$, $b_2 = a_3$, $b_3 = a_5$, and so on. In general: $b_{4\mathrm{n}} = a_{2\mathrm{n}}$; $b_{2\mathrm{n}+1} = a_{4\mathrm{n}+1}$; $b_{4\mathrm{n}+2} = a_{4\mathrm{n}+3}$.
Let ODD-NEW $\equiv$ selected individual's new label b_{n} is odd
EVEN-NEW $\equiv$ selected individual's new label b_{n} is even

5. Since the new labelling is just as good as the old one, the reasoning in step 2 ought to show that $Pr_{\mathrm{R}}(\text{ODD-NEW}) = Pr_{\mathrm{R}}(\text{EVEN-NEW}) = \frac{1}{2}$.

But ODD-NEW represents the same set of individuals as ONE, so that we must also have $Pr_{\mathrm{R}}(\text{ODD-NEW}) = Pr_{\mathrm{R}}(\mathrm{ONE}) = \frac{1}{4}$,

a contradiction.

The key assumptions in deriving the contradiction are equiprobability and *Label Independence*: re-labelling (or re-arranging) the members of a countably infinite population should make no difference to probability claims. It is tempting to see the equiprobability assumption as the cause of trouble, but in my view, *Label Independence* is the culprit. Re-labelling does make a difference.

To defend this claim, we need to become more systematic about the basis for assigning relative betting quotients, and in particular RBQ = 1, which reflects a judgement of equiprobability. In many cases where directly relevant experience is unavailable, and certainly in abstract mathematical examples such as the de Finetti lottery and the present paradox, we base our judgements on intuitions of symmetry. If, so far as all knowledge that could influence our assessment of likelihood goes, there is no basis for distinguishing between two sets of outcomes (i.e., we lack knowledge of any asymmetry), then we are inclined to regard them as equiprobable. We feel confident in asserting equiprobability even when (as in the de Finetti case) we are unsure about whether we can assign a probability value.

Obviously, we are appealing here to a version of the Principle of Indifference (replacing 'equal probability' with 'equiprobability' in the traditional formulation). Given the paradoxes that arise from reckless use of the Principle of Indifference, however, a *necessary* condition for its application is that the relevant 'symmetries' do not lead to incoherent claims—as has just happened with our paradox.

Let us say, by extension, that a set of symmetries is *coherent* if it passes this test. I believe that coherence and having a well-motivated story about the underlying symmetries are *sufficient* to justify assertions about relative probabilities, but I will not defend that claim here. It is worth noting, though, that worries about incoherence have traditionally been the major objection to the Principle of Indifference. In any case, the view I am urging is that, in many cases, the problem of determining whether a set of equiprobability judgements is coherent reduces to showing that the underlying set of symmetry judgements is coherent.

The next section outlines a conception of symmetries sufficient for defining the relative probabilities (or relative betting quotients) in the rather abstract settings that have occupied our attention so far. In particular, we shall see one necessary condition for coherence. This condition provides us with a principled way to distinguish between the plausible way we have applied symmetry-based arguments to the de Finetti lottery and the reckless way we applied them to the re-labelling paradox. The claims about equiprobability in the de Finetti lottery can be justified from the underlying symmetries, but the claims made for the re-labelling example in the present section cannot.

5.2 Resolution: from symmetry to relative probability

Following earlier work (Bartha and Johns [2001]), symmetries are construed as bijections on a space X of outcomes, with certain additional properties that depend upon features of the space. A symmetry is meant to be a mapping that preserves all features that bear on probability, and these features vary from case to case.

Definition 1 (Regularity): A set S of symmetries is *regular* if it has the following two properties:

(G) S is a group under function composition

(M) For no non-empty subset C of X and positive integers $m > n$ are there symmetries $\theta_1, \ldots, \theta_m$ and $\Psi_1, \ldots, \Psi_n$ in S such that

$$\bigcup_{i=1}^{m} \theta_i(C) \subseteq \bigcup_{j=1}^{n} \Psi_j(C)$$

where the union on the left is disjoint.

One trivial example of a regular set of symmetries is the group of all permutations on the set $\{1, \ldots, n\}$ (i.e., the symmetry group of order n).

The basic idea of using symmetries to define a relation of relative probability is as follows. If $B = \theta(A)$ for a symmetry mapping θ, we will want to say that B and A are equiprobable and that the relative probability $R(B, A) = 1$. If $B = \theta_1(A) \cup \theta_2(A)$ for distinct symmetries θ_1 and θ_2, and the union is disjoint, we want to say $R(B, A) = 2$. More generally, $R(B, A) = n$ if B can be written as the disjoint union of n copies of A under symmetry mappings. This can be extended to non-integer values, as described in (Bartha and Johns [2001]) and summarized briefly in the Appendix to this paper. The problem, of course, is to make sure that these relative probabilities are well-defined.

The condition (M) rules out the possibility that m disjoint copies of C could be placed inside n copies of C even though $m > n$. This means, for example, that C cannot be written as the disjoint union of two or more copies of itself under symmetry mappings. It should be clear that this had better not happen if we want to use S to define relative probabilities, because in such a case we would have to say that C is both equiprobable with, and twice as probable as, itself! Bartha and Johns ([2001]) argue that both (G) and (M) are necessary conditions for a set of symmetries S to induce a coherent relationship of equiprobability and, more generally, relative probability on pairs of sets of outcomes.

In the case of the de Finetti lottery, we start by identifying the outcome space as the set of all positive integers. We have already made implicit appeal to two types of symmetry transformations. First, in Section **3.2**, we suggested that the even and odd tickets might each have relative probability 0.5 compared to the set of all tickets; this rests on the idea that the translation $\sigma(x) = x + 1$, and more generally translations by a fixed integer, $\sigma(x) = x + k$, count as symmetries. Any two sets of tickets related in this way are equiprobable. Second, in Section **3.1**, we suggested that a swap of any two tickets should not affect any of our probabilities, and we argued in Sections **2** and **3** that 'ticket i wins' and 'ticket j wins' should be equiprobable. These claims rest on the idea that any permutation τ of finitely many integers counts as a symmetry. The Appendix provides a proof that the set S consisting of products of mappings of these two types is in fact coherent. So this set of symmetries induces well-defined relations of relative probability and equiprobability. This provides justification for the fair-betting quotient approach of Section **4**.

By contrast, things go wrong when we turn to the paradox of Section **5.1**. The outcome space is effectively the same: the set of all positive integers. The intended set of symmetries, however, is much larger. In fact, *Label Independence* is equivalent to the supposition that *any* bijection on the natural numbers (which is what a re-labelling amounts to) is an acceptable symmetry transformation that can be used to justify an assertion of relative probability. But there is a bijection that takes ODD to ONE, and (obviously) another that takes ODD to THREE. Since ODD = ONE ∪ THREE, we are in trouble: we have violated the condition (M), and our set of symmetries is incoherent.

We have to reject *Label Independence*. We cannot allow just any permutation, or re-labelling, of our population and hope to preserve probabilistic relationships. There is simply no way to define a label-independent uniform (relative) probability distribution on a countably infinite population. The condition (M), then, provides a useful test to rule out inappropriate appeals to symmetry.

We are left with the problem of explaining the original intuition that *Label Independence* is reasonable for a countably infinite set. Why does it matter which labels we attach to individuals? In order to answer this question, consider two closely analogous cases.

Example 2: In any finite set of outcomes, *Label Independence* is valid. Re-labelling the elements does not disturb an initial uniform probability distribution.

Example 3: In the outcome space of real numbers in the interval [0, 1], *Label Independence* is not valid. An initial uniform probability distribution assigns probability $\frac{1}{2}$ to the intervals $[0, \frac{1}{2}]$ and $[\frac{1}{2}, 1]$. We can re-label by compressing $[0, \frac{1}{2}]$ into $[0, \frac{1}{4}]$, and expanding $[\frac{1}{2}, 1]$ into $[\frac{1}{4}, 1]$. That is, define

$$f(x) = \begin{cases} \frac{1}{2}x, & 0 \leq x \leq \frac{1}{2} \\ \frac{3}{2}x - \frac{1}{2}, & \frac{1}{2} \leq x \leq 1 \end{cases}$$

and re-label each point x as $f(x)$. If re-labelling did not matter, we would have to conclude that $Pr_{\mathrm{R}}([0, \frac{1}{2}]) = Pr_{\mathrm{R}}([0, \frac{1}{4}])$, which is false.

Because of the nature of the usual probability measure on [0, 1], metric relationships, or inter-point distances, are relevant to probability assignments. The mapping discussed in Example 3 is not an acceptable symmetry because these metric relationships are not preserved—though a mapping that permuted finitely many points and left everything else undisturbed would still be fine. We now see that a similar point applies to a countably infinite set A.

Mappings that distort inter-point distances for more than a finite set of points are not acceptable symmetries. A probability measure on A, or more generally a relationship of relative probability, depends not just on the population A but also on a specific ordering or arrangement of that population.

Comparing the three cases—a finite outcome space, a countably infinite outcome space, and the space [0, 1]—it is puzzling that a uniform probability distribution is unproblematic and can be reconciled with countable additivity in a finite space and for the interval [0, 1], but there is apparent incompatibility for a countably infinite space. Symmetry-based relative probabilities restore the analogy by providing a common framework for understanding all three cases: each time, an underlying set S of symmetries is responsible for the fundamental relationships of uniformity or equiprobability.

6 Beyond the de Finetti lottery

We began with two contrary positions. The Dutch Book argument purports to show that if we ground subjective probabilities in a betting formalism, countable additivity is inevitable. Kelly argues that the axiom is plainly inapplicable to the de Finetti lottery, and more generally too powerful to be taken for granted in subjective approaches to the confirmation of scientific theories.

The foregoing analysis removes one bone of contention: the de Finetti lottery. We can still work within a betting framework while accepting that countable additivity does not apply to cases such as the de Finetti lottery, where we have good reason to limit ourselves to relative probabilities. But ought we to be worried about countable additivity outside of such cases? In particular, ought we to be worried about its use in the confirmation of scientific hypotheses? I want briefly to argue that there is no real basis for extending the lessons of the de Finetti lottery to full-blown skepticism about the axiom (assuming, of course, that we are still not calling Dutch Book arguments into question).

Let us focus, as Kelly does, on the problem of confirming a universal hypothesis through repeated trials. Countable additivity entails that if such a hypothesis is false, it will be refuted sooner rather than later. Kelly ([1996], p. 324) acknowledges that if we have a sequence of data known to come from independent and identically distributed (i.i.d.) trials, the bias towards early counterexamples is unproblematic. In consequence, his reservations about countable additivity are restricted to contexts where we lack this knowledge of independence, as illustrated by the following example (Kelly [1996], p. 323).

Example 4: Sextus is given a sequence of observation reports, each of which is simply a 1 or a 0. The data set at any given time is thus a finite string of 1s and

0s. Consider the hypothesis T that every observation report will be a 1. Sextus points out that no matter how many 1s have been reported, the next report might be a 0.
Let H_n stand for the proposition that the first 0 appears as the nth report. Countable additivity ensures that most of the probability that T is false is distributed over some initial group of H_ns. Provided that we start with a non-zero prior probability for T, Sextus can come as close as we please to a posterior probability of 1 after a sufficiently long string of 1s.

It is precisely in this sort of case, where we lack knowledge of whether the data are generated by independent trials, that Kelly is unhappy with (CA) and its accompanying bias towards early refutation. Why not follow de Finetti, and assign each H_n probability 0 (rejecting countable additivity)? Or why not follow the approach to the de Finetti lottery outlined above and insist that $R(H_n, H_m) = 1$ for all n, m?

Upon first consideration, it looks as though the symmetry arguments employed in the de Finetti lottery can be extended to this example. If 'All observations will be 1s' is false, then we will see at least one counterexample in a countably infinite run of trials, and exactly one H_n will be true. Lacking knowledge of independence, any two trials appear to be symmetrical with regard to producing the first counterexample, and so the propositions H_n should be equiprobable. As in the case of the de Finetti lottery, we should distribute the probability of generating the first counterexample uniformly over the countable infinity of trials. But that is just what countable additivity forbids.

On closer examination, however, this argument is inconclusive. One disanalogy between this example and the de Finetti lottery, of course, is the temporal ordering of the trials. A second important disanalogy is that in the case of the de Finetti lottery, we know that one ticket will win, but in the case of Sextus, we do not know whether a 0 will eventually show up. Together, these disanalogies undermine the symmetry argument of the preceding paragraph. What the symmetries present in this situation establish is only that a result of 1 on observation m is equiprobable with a result of 1 on observation n, for any m and n (since any two observations are alike, for all Sextus knows). If we knew that a single 0 would appear in an infinite run, then the situation would become directly analogous to the de Finetti lottery, and we would have a justification for maintaining that H_n and H_m are equiprobable. But we don't know this. Hence, the time asymmetry cannot be ignored (since H_n entails that nothing but 1s appear for the first $n - 1$ trials), and there is no basis for the assertion of equiprobability.

To assume a uniform prior over the propositions H_n is to assume with probability 'nearly' 1 that *no* early counterexample will be found. Confirmation for T by any finite set of observed 1s becomes impossible, no matter how large the data set. This is not justified by the symmetries in the set-up. As far as Sextus is aware, the stream of 1s might be generated in any number of ways, including the following:

(a) a coin is tossed repeatedly, with a report of 1 on heads and 0 on tails
(b) a de Finetti lottery is secretly conducted, and 1 is reported for losing ticket numbers and 0 for the single winning ticket
(c) somebody writes down 1s *ad infinitum*.

To assume a uniform distribution over the hypotheses H_n is to assume that the method employed to generate the data is like that of the second or third suggestion, ruling out (for example) the idea that they come from independent trials.

It might be objected that we can make use of arguments similar to those discussed in Section **3.1**. For every mechanism that favours H_n over H_m, there is another that simply swaps the results of trials *m* and *n* before generating the data sequence. Is it not then permissible to adopt a uniform probability distribution? This objection fails because each H_n contains information not just about trial *n*, but also about all the preceding trials. Nor does it seem that we can find any coherent, well-motivated set of symmetries that would justify treating H_n and H_m on par. H_n and H_m represent initial data sequences of different lengths. Additional knowledge (as in the de Finetti lottery) would be needed to generate a relationship of symmetry.

In short, it is certainly possible for Sextus to adopt a uniform distribution over the hypotheses H_n, by refusing to assign any betting quotients and insisting that the trials are generated by a method like (b) or (c). It is not possible, however, to claim that this uniform distribution is well-motivated and justified by the symmetry of the situation, as is the case for the de Finetti lottery. For Sextus, and even more for real-world testing of scientific hypotheses, the temporal ordering of tests and the possibility of independence undercut the symmetry arguments.

For a symmetry-based justification of a uniform probability distribution over a countable partition, three things are needed:

1. a rationale for refusing to assign betting quotients to the elements of the partition; ideally, as in the de Finetti lottery, this means an argument that *no* real-valued betting quotient can be fair
2. an epistemic situation in which any two of these elements are perfectly symmetrical so far as probabilistic considerations go, so that any permutation that swaps just two elements counts as a symmetry
3. an underlying set of symmetries that is coherent.

Absent these conditions, it is not clear that we can make sense of the idea that countable additivity fails but the other standard probability axioms remain. A judicious use of symmetry arguments supports dropping countable additivity in the case of the de Finetti lottery, but defuses Kelly's threat to more ordinary applications of the axiom.

Acknowledgements

For extremely helpful discussions, criticisms and suggestions, I am indebted to Richard Johns, Christopher Hitchcock, Peter Vranas and two anonymous referees.

Department of Philosophy
1866 Main Mall, E-370
University of British Columbia
Vancouver, B.C.
V6T 1Z1 Canada
bartha@interchange.ubc.ca

Appendix

This Appendix extends the definition of the relative probability function to rational values, and verifies the claim of Section **5.2** that the set of symmetries that pertain to the de Finetti lottery generates a well-defined relative probability function (or, what amounts to the same thing, well-defined relative betting quotients).

Definition 5 (Relative Probability): Suppose X is a set of outcomes and S is a set of symmetries on X. Two subsets A and B are said to be commensurable if either of the following two conditions is met, with the relative probability defined accordingly:

1. $A = \bigcup_{i=1}^{m} \theta_i(C)$ and $B = \bigcup_{j=1}^{n} \Psi_j(C)$ for some C, with both unions disjoint. In this case, set $\mathrm{R}(A, B) = m/n$. (C is called a *divisor* of A and B.)

2. Both $\mathrm{R}(A, C)$ and $\mathrm{R}(B, C)$ are defined and $\mathrm{R}(B, C) \neq 0$ for some C. In this case, set $\mathrm{R}(A, B) = \mathrm{R}(A, C)/\mathrm{R}(B, C)$.

In general, we have no assurance that $\mathrm{R}(A, B)$ is well-defined. The following results extend ideas developed in Bartha and Johns ([2001]). That paper shows that a commutative group of symmetries yields a well-defined relative probability function. The group of symmetries we want to apply to the de Finetti lottery situation, however, is not commutative, for it includes both shifts by a fixed integer and permutations of finitely many integers. It is not difficult to see, however, that this group has a property I shall call *-commutativity. First, we shall see that R is well defined for *-commutative symmetry groups; then, we shall return to the de Finetti lottery.

Definition 6 (*-commutativity): A set of symmetries S is **-commutative* if for any θ, θ' in S, there is a finite (possibly empty) subset F of X such that $\theta\theta'(x) = \theta'\theta(x)$ for all $x \in X \setminus F$. (The two symmetries commute outside of the set F.)

Lemma 7: Suppose S is a *-commutative group of symmetries on X. Suppose that for some infinite subset C of X,

$$\bigcup_{i=1}^{m} \theta_i(C) \subseteq \bigcup_{j=1}^{n} \Psi_j(C) \cup K$$

where the union on the left is disjoint and K is finite. Then $m \leq n$.

Proof: First, let F be a finite subset of X that includes all the finite subsets on which any two mappings in $\{\theta_1, \ldots, \theta_m, \Psi_1, \ldots, \Psi_n\}$ fail to commute; indeed, take F large enough so that outside F, up to three applications of these functions or their inverses commute. For example, if $x \notin F$, then $\theta_1\Psi_2^{-1}\theta_2(x) = \Psi_2^{-1}\theta_2\theta_1(x)$. Let N be the number of elements in K.

Suppose, for a contradiction, that $m > n$. Choose $N + m$ elements of C all lying outside F. For at least one of these elements, c, it must be that $\theta_i(c) \in \Psi_j(c)$ and $\theta_{i'}(c) \in \Psi_j(c)$ for the same j but $i \neq i'$. Without loss of generality, we may assume $\theta_1(c)$ and $\theta_2(c)$ are both in $\Psi_1(c)$. But then $\Psi_1^{-1}\theta_1(c) \in C$ and so $\theta_2\Psi_1^{-1}\theta_1(c) \in \theta_2(c)$, and similarly $\theta_1\Psi_1^{-1}\theta_2(c) \in \theta_1(c)$. Since c is not in F, everything commutes and $\theta_2\Psi_1^{-1}\theta_1(c) = \theta_1\Psi_1^{-1}\theta_2(c)$. But then $\theta_1(c) \cap \theta_2(c)$ is not empty, contradicting our assumption.

Proposition 8: If S is a *-commutative group of symmetries, then $R(A, B)$ is well-defined (i.e., whenever some value can be assigned, this value is unique).

Proof: First suppose that A and B have two distinct divisors, C and D. Write $A = mC$ and $B = nC$ to indicate that A can be written as m copies of C and B as n copies of C, as in the definition. Suppose also $A = m'D$ and $B = n'D$, where the relevant symmetries are as in the diagram. We have to show that $m/n = m'/n'$. There is no problem at all if C is a finite set (in which case A, B and D must all be finite): counting elements suffices to establish the result. So we may assume that C is an infinite set.

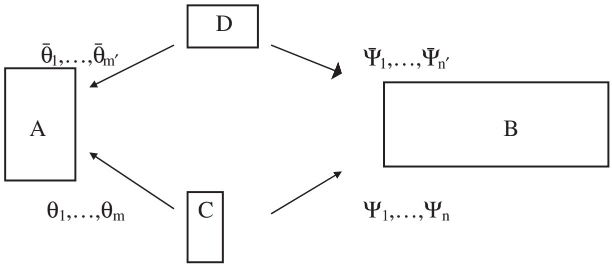

For each $c \in C$, and for $1 \leq i \leq m$, $\theta_{\mathrm{i}}(c) = \overline{\theta}_{\mathrm{k}}(d)$ for some $d \in D$ and some k with $1 \leq k \leq m'$; also, for $1 \leq j \leq n'$, $\overline{\Psi}_{\mathrm{j}}(d) = \Psi_{\mathrm{l}}(c^*)$, for some $c^* \in C$ and some l with $1 \leq l \leq n$. This shows that for any $c \in C$ and any subscripts i and j, there are $c^* \in C$ and subscripts k and l such that

$$\overline{\Psi}_{\mathrm{j}}(\overline{\theta}_{\mathrm{k}}^{-1}(\theta_{\mathrm{i}}(c)) = \Psi_{\mathrm{l}}(c^*)$$

Write $C = C_1 \cup C_\infty$, where C_1 is a finite set such that any three of the symmetry mappings or their inverses commute outside C_1, and C_∞, which includes everything else in C, is infinite. This means that if $c \in C_\infty$, it follows that $\overline{\Psi}_{\mathrm{j}}\theta_{\mathrm{i}}(c) = \overline{\theta}_{\mathrm{k}}\Psi_{\mathrm{l}}(c^*)$ (using *-commutativity), where $c^* \in C$. Now c^* might be in C_1 or C_∞, so we have

$$\bigcup_{i=1}^{m}\bigcup_{j=1}^{n'}\bar{\Psi}_j\theta_i(C_\infty) \subseteq \bigcup_{k=1}^{m'}\bigcup_{l=1}^{n}\bar{\theta}_k\Psi_l(C) \subseteq \bigcup_{k=1}^{m'}\bigcup_{l=1}^{n}\bar{\theta}_k\Psi_l(C_\infty)\bigcup K,$$

where K is finite. So by Lemma 7, $mn' \leq m'n$. But analogous reasoning gives a similar containment in the reverse direction, so $mn' = m'n$.

For the general case, A and B are linked by intermediates C and D, as in clause (2) of the above definition of relative probability. The picture is the same as in the diagram, except that (following the method of Bartha and Johns [2001]) we link each of A and B to each of C and D by a finite sequence of intermediate sets, such that each pair in succession has a divisor. The resulting diagram has two sets of zigzag patterns joining A to B, one via C and the other via D. By constructing the same sort of argument for C as just given, *-commutativity and Lemma 7 establish that R(A, C)/R(B, C) = R(A, D)/R(B, D).

There are two ways to apply these results to the de Finetti lottery. The first approach identifies the outcome space as $\{1, 2, 3, \ldots\}$, i.e., the set of all positive integers. In this case, the symmetry group S must consist solely of permutations τ of finitely many integers. Shifts of the form $\theta(x) = x + d$, where d is a positive constant, have no inverses and, indeed, are not bijections. This gives us a *-commutative group of symmetry transformations, so Proposition 8 shows that we get well-defined relative probabilities. Further, for each pair i and j of positive integers, there is a permutation τ in S that swaps i and j but does nothing else, which establishes that any set of outcomes including i is equiprobable with the same set, substituting j for i. This is enough to vindicate the idea that 'ticket i wins' and 'ticket j wins' are equiprobable, provided we accept these underlying symmetries as well-motivated.

We need a slightly more elaborate approach to justify the intuition that holding all even-numbered tickets is no better or worse than holding all odd-numbered tickets. This claim can be defended if we are able to include the symmetry mapping $\theta(x) = x + 1$ in our set of symmetries S. Start with the outcome space $\{0, \pm 1, \pm 2, \ldots\}$ of all integers, and let the set of symmetries S consist of all finite products of shifts $\theta(x) = x + d$ and finite permutations τ. In fact, each member of S is of the form $\sigma = \theta\tau$, but the only point that concerns us is that S is a *-commutative group. Hence, we obtain a well-defined relative probability function on the integers. To apply this to the de Finetti lottery, we can either extend the lottery by including negative-numbered tickets, which makes no essential difference to the debate, or we can restrict the relative probability function just obtained to subsets of the positive integers.

References

Bartha, P. and Hitchcock, C. [1999]: 'The Shooting-Room Paradox and Conditionalizing on "Measurably Challenged" Sets', *Synthese*, **118**, pp. 403–37.

Bartha, P. and Johns, R. [2001]: 'Probability and Symmetry', *Philosophy of Science (Proceedings)*, **68**, pp. S109–22.

de Finetti, B. [1972]: *Probability, Induction and Statistics*, New York: Wiley.

de Finetti, B. [1974]: *Theory of Probability*, Vols 1 and 2, trans. A. Machí and A. Smith, New York: Wiley.

Edwards, W., Lindman, H. and Savage, L. J. [1963]: 'Bayesian Statistical Inference for Psychological Research', *Psychological Review*, **70**, pp. 193–242.

Howson, C. and Urbach, P. [1993]: *Scientific Reasoning: The Bayesian Approach*, 2nd edn, La Salle: Open Court Press.

Jaynes, E. T. [1985]: 'Some Random Observations', *Synthese*, **63**, pp. 115–38.

Jeffrey, R. [1957]: 'Contribution to the theory of inductive probability', Ph.D. dissertation, Princeton University.

Jeffreys, H. [1961]: *Theory of Probability*, Oxford: Oxford University Press.

Kelly, K. [1996]: *The Logic of Reliable Inquiry*, Oxford: Oxford University Press.

Lewis, D. [1980]: 'A Subjectivist's Guide to Objective Chance', in R. Jeffrey (*ed.*), 1980, *Studies in Inductive Logic and Probability*, Berkeley: University of California Press, pp. 263–93.

Maher, P. [1997]: 'Depragmatized Dutch Book Arguments', *Philosophy of Science*, **64**, pp. 291–305.

McCall, S. and Armstrong, D. [1989]: 'God's Lottery', *Analysis*, **49**, pp. 223–4.

Savage, L. [1972]: *The Foundations of Statistics*, New York: Dover.

Skyrms, B. [1984]: *Pragmatics and Empiricism*, New Haven: Yale University Press.

Spielman, S. [1977]: 'Physical Probability and Bayesian Statistics', *Synthese*, **36**, pp. 236–69.

Williamson, J. [1999]: 'Countable Additivity and Subjective Probability', *British Journal for the Philosophy of Science*, **50**, 401–16.