An Outcome of the de Finetti Infinite Lottery is Not Finite

Marc Burock

Abstract

A randomly selected number from the infinite set of positive integers—the so-called de Finetti lottery—will not be a finite number. I argue that it is still possible to conceive of an infinite lottery, but that an individual lottery outcome is knowledge about set-membership and not element identification. Unexpectedly, it appears that a uniform distribution over a countably infinite set has much in common with a continuous probability density over an uncountably infinite set.

Introduction

De Finetti [1974] wanted to know if we could make sense of a random lottery containing countably infinite tickets using standard axiomatic probability theory. He discovered a problem. Given a countably infinite set of options, it appears impossible to assign any non-zero probability to each option without violating the basic axioms. If the lottery contained only one-billion tickets there would be no problem—simply assign one billionth to each ticket. But infinity is the number we are dealing with, and one divided by infinity is zero according to our mathematicians. Although assigning zero probability to each ticket is not a problem in itself, once we assume that at least some ticket is selected, it is contradictory to say that the probability of that selection was exactly zero.

There have been attempts to resolve the contradiction implied by the de Finetti lottery, for instance, by abandoning the troublesome axiom that yields the contradiction or invoking non-standard measure theory (Bartha [2004], Vallentyne [2000]). Having strong empiricist tendencies, I was compelled to simulate an approximate infinite lottery on a computer. Of course it is impossible to simulate countably infinite options on a finite computer. In fact, this has been an argument itself against the infinite de Finetti lottery— no physical mechanism can make selections over an infinite range of elements (Howson and Urbach [1993]). Fortunately one need only do an approximate simulation to get a feel for what the true infinite lottery would look like.

burock@alum.mit.edu

Simulated lotteries

Do this on a computer. Start with a set of ten numbers and perform an equally likely (random) selection across them. No real surprises here. On average you see each number, one through ten, appear after performing a dozen or so simulations. Now increase your lottery to have one billion tickets. There is now a 99% chance that you will see a number greater than ten-million on the first simulated outcome. While all numbers are equally likely, you are almost certain to see a number that *looks* large because large looking numbers are so common[1]. Extending the range of the lottery even further, perhaps to a number represented by a one followed by one-thousand zeros, you are nearly guaranteed to select a seemingly gigantic number even if you make one billion selections. In this case you are nearly certain to choose a number with *hundreds* of digits. That seems pretty big. As the number of tickets in our lottery tends to infinity, it would appear that we are sure to choose a really, really, huge number; a number bigger than any finite number – infinity.

Lottery analysis

A simple limit analysis will clarify the situation somewhat and show that the simulation is unneeded to justify the conclusion. Suppose we have a lottery with N tickets that are labeled one through N. Now choose any finite number $a<N$. Let x represent the value of the ticket that is randomly selected where each ticket is equiprobable, then

$$P(x \le a) = a / N$$
$$P(x > a) = (N - a) / N$$

It is easy enough to evaluate the limit as N goes to infinity (zero for the top expression and one for the bottom). Regardless of the value you specify for a, so long as a is finite you are assured with probability one that the value obtained in the infinite lottery will be greater than a. In other words, if you were to bet on an infinite lottery, simply ask your competitor to write down the largest number she knows, then bet that the number drawn from the lottery will be bigger than the number she specified and you are certain to win. Further, since the outcome of the lottery will be greater than a for any assigned value a, the outcome x cannot be finite, as this would imply that we could assign a value to a that was finite and greater than x (such as $x+1$), which would contradict the assumption that $P(x \le a)=0$.

[1] While ten-million may appear to be a large number, on a scale between zero and one-billion it is small. Not surprisingly, it is among the one-percent smallest numbers on that scale.

The above argument is consistent with the inference drawn from the pseudo-simulations: an outcome of the de Finetti infinite lottery is not finite. To make this argument I have assumed

1) Equiprobability between outcomes (Principle of Indifference)
2) Non-zero probability for each outcome before taking the limit

I cannot envision how one would physically run an infinite lottery, nor am I certain what an outcome of an infinite lottery would look like, but I am fairly sure that it will not be a finite ticket number and I will place good money against anyone who says otherwise. It appears that an outcome of the de Finetti lottery will be larger than any a priori conceived finite value, which is not too troubling given that we are working with an infinite number of alternatives. Can we say that the outcome of this lottery is infinity? Perhaps, but the meaning of that statement is not entirely clear. For instance, it does not imply that there exists a ticket with the infinity symbol drawn on it, and that this ticket will be selected with certainty. The infinite ticket does not exist as such. We can only imagine this infinite ticket because of the order-type of the natural numbers.[2] Changing the ordering may change the result. A lottery over the integers using a similar analysis as above would also suggest a non-finite outcome, however, now both positive and negative infinity become potential outcomes. Can we assign ½ probability to each?

Let us first consider an analogous example that avoids order. Instead of the well-ordering of the de Finetti lottery over the natural numbers, assume that there is a countably infinite but unordered set, for instance, an infinite stack of playing cards with a unique picture on each card. Without ordering it makes little sense to speak of the infinite card. Now arbitrarily partition the stack of cards into a *finite* subset of cards and an *infinite* subset of cards using any mechanism you like, physical or not. Keep in mind that the notion of countable infinity, as it is commonly understood in the Cantorian sense, allows one to partition an infinite set into a finite set and another infinite set with the same cardinality (size) as the original set. Next ask someone to choose at random a card from the *entire* collection of cards—the person chooses not knowing that you have made an arbitrary finite partition, and your partition does not affect the selection process. Since the infinite partition of cards is infinitely larger than the finite partition, and the person is selecting a card at random, the odds and probability that the selected card will be in the finite partition are zero.[3] It follows that the selected card will be a member of the infinite partition. Further, since the outcome will be a member of the infinite partition for any given finite partition, the outcome cannot be contained in *any* finite partition a priori, as

[2] I use the word 'order' in the technical, set-theoretic sense. The natural numbers and integers are both countably infinite sets that have the same cardinality but differ in order-type.
[3] The limit analysis here is identical to the analysis used above, where 'the size of the finite set' in this example takes the place of the arbitrarily chosen finite number *a* in the previous example.

this would imply that we could create an arbitrary finite partition that contained the selected card, which would contradict the assumption that the probability of finding the card in the finite partition is zero.

This conclusion can be generalized to the de Finetti infinite lottery without altercation, and forms the main assertion of this paper: Given a uniform probability distribution over a countably infinite set the outcome of a selection will not be a member of any a priori identified finite set of alternatives. It seems reasonable to interpret the a priori impossibility of finite set membership as a zero probability measure over every conceivable finite set of alternatives. Although the outcome will not be a member of any finite set of alternatives, we conclude without contradiction that the outcome will be a member of a countably infinite set of alternatives. Since non-zero probability measure is only defined over an infinite set taken as a single object – and this is important – the 'outcome' of an infinite lottery can only tell us the set-membership of the selected element and nothing about the particular element itself. Why is this so? I argue out of consistency. If all finite sets of alternatives have zero probability measure, and only infinite sets have non-zero probability, then it is inconsistent to acknowledge a particular a posterior finite outcome. Consider again the de Finetti lottery. It seems clear that the drawn ticket will not be any finite number. We may say that the outcome will be infinitely large, but that too is ambiguous and depends upon the order-type of the natural numbers. It is clearer and more general to say that the outcome of the de Finetti lottery will be a member of an infinite set of alternatives—for natural numbers that implies an infinite set 'off to the right'. While this statement appears to be a tautology (it is an assumption of the de Finetti lottery), it does provide new information. Most pertinently, you would never bet that the outcome will be found in any finite set of a priori alternatives.

The 'outcome' of an infinite lottery

Although until this point I have not made a direct appeal to infinitesimal probabilities, it seems clear that I have assumed the existence of infinitely small probability measures that when collected into an infinite-sized set sum to a finite number. Non-standard infinitesimal probabilities have been used to make sense of countably infinite sets in previous work (Vallentyne [2000], Bartha and Hitchcock [1999]), however, the conclusions and application in the present work differ considerably. While I acknowledge that the outcomes in the infinite lottery are equiprobable, I deny that our outcome will be contained in any finite set of alternatives. If my knowledge of real-analysis were adequate, I believe I could say in a more rigorous manner—through formal proof—that the outcome of an infinite lottery will *almost surely* not be contained in any finite set of alternatives. Further, since non-infinitesimal probability measure can only be defined over a countably infinite set taken as whole where each individual alternative is almost surely not going to occur, an outcome corresponding to this measure is one of set-membership and not of element identification. In other words, for an infinite lottery one cannot know the ticket that was chosen, but only the set to which the ticket belongs. That is a strange situation, but it seems to be 'empirically' supported by inference from a large but finite lottery.

Our apparently strange conclusion here is not completely novel, in fact, Bartha [2004] in an analysis of the de Finetti lottery using the concept of relative betting quotients came to a similar conclusion stating that "…the propositions 'ticket *n* wins' and 'some ticket wins' are probabilistically incommensurable." We are told these propositions are incommensurable because the relative betting quotients for 'ticket *n* wins' are undefined, although it would not be too great a leap to say that the probability that 'ticket *n* wins' is infinitesimal while 'some ticket wins' is a finite, and therefore we cannot directly compare these propositions. In light of the preceding analysis, I am not completely confident that the proposition 'some ticket wins' even makes sense. The best we can say in the language of probability is that 'the chosen ticket belongs to an infinite set'. Selection of a ticket in an infinite lottery is an almost imaginary process that tells us nothing about the ticket other than the set it belongs too. I cannot help but feel that the Axiom of Choice has an important part to play in the random selection of one element from an unordered countably infinite set. We cannot constructive demonstrate how such a selection would take place, yet we intuit that it is possible in some sense, relying upon an axiom to get the job done. Perhaps the price we pay for invoking the Axiom of Choice is incomplete knowledge about the actual selection.

A 'mixed' discrete distribution and the negated event

In the following example I will show that uniform distributions over countably infinite sets may arise from a negated event. Additionally, negation is coupled to the notion that the information of a probabilistic outcome may be limited to set-membership information and nothing more. The example will make these words more obvious.

Example 1. A 'mixed' distribution and the negated event. Suppose the only thing we know about a probability distribution over the natural numbers is that the probability of selecting the number one is 1/4. It follows that the probability that one is *not* selected is given by $P(x = \neg 1) = 3/4$.

Since we are told that the sample space is the set of natural numbers **N**, the event $x = \neg 1$ is equivalent to the event $x \in \{2, 3, 4, ...\} = \neg 1$. We may invoke a symmetry argument (indifference argument) and reason that the probability 3/4 is uniformly spread over the countably infinite collection of elements in {2, 3, 4,…}. The 'mixed' distribution is explicitly:

$$P(x=1)=1/4$$
$$P(x=\neg 1)=P(x\in\{2,3,4,...\})=3/4$$

I call this distribution a mixed distribution because it is analogous to so-called mixed probability density functions where part of the density function is infinitely concentrated at discrete locations and part is finitely concentrated but continuously spread. In this example the probability is 'infinitely dense' at x=1 relative to other values. Anyone who has worked with mixed density functions will find this analogy compelling. It appears that de Finetti's infinite lottery is not completely contrived, and is somehow analogous to the probability of a negated event that is contained in an infinite sample space.

The idea that the outcome of a uniform infinite selection is one of set-membership and not element identification directly follows from this example of the negated event, however, any negation in a finite example would have yielded a similar result. Here is another hand-waving explanation of why all outcomes of uniform infinite selections correspond to set-membership. Given a uniform infinite selection, we know that the outcome will be non-finite, and perhaps this necessary negation implies that we cannot know the selected element but only the membership of that element. It has been argued elsewhere that the negated event plays a special role in connecting with the sample space in probability theory (Burock [2005]).

Density functions and uniform countably infinite distributions

Continuous density functions over the real line **R** can help us understand the mysterious de Finetti infinite lottery. We rarely acknowledge in probabilistic practice that the sets of outcomes on **R** are typically uncountably infinite. When manipulating continuous density functions and computing probability, the question almost always takes the following form: what is the probability that outcome x falls within the interval [a, b], or $P(a\leq x\leq b)$? The number of elements contained between the interval [a, b] can be denoted by the infinite cardinal number $\aleph_1$. The analogy to the de Finetti infinite lottery should be clear. I have argued that for a uniform distribution over a countably infinite set N, it only makes sense to ask about $P(x\in A)$ where A is some countably infinite subset of N, just like over an uncountably infinite set R, it only makes sense to ask $P(x\in[a,b])$ where [a, b] is some uncountably infinite subset of R. The probability of each element in A is zero or infinitesimal, just like the probability of each element of [a, b]. An obvious disanalogy is that the density over [a, b] need not be uniform. In the Appendix of this work I show how the concept of density can be applied to countably infinite sets.

Conclusions

After all of this one may still counter quite easily that the de Finetti lottery is simply impossible. Perhaps, but for consistency one should then also declare all continuous probability density functions impossible as well; a situation that would deprive quantum physics of its most useful tool. One may view the de Finetti countably infinite lottery as a bridging concept between finite probability examples and continuous examples. More mathematically capable individuals should be able to unify the two more rigorously than was done here. The outcomes of uniform countably infinite draws can only give us information about set-membership, analogous to continuous probability on the real line. Importantly, not all uniform countably infinite distributions are equivalent. The formal set-theoretic order-type of an infinite set will differentiate one uniform countably infinite distribution from another.

Appendix

The Density of a Countably Infinite Set

This Appendix suggests how the concept of density may be extended to countably infinite sets, and provides several examples illustrating how this concept may be applied to answer probabilistic questions. Nothing below (or in this entire work) is mathematically rigorous.

Here is the situation I wish to understand. Suppose we are given the set $\mathbf{N}=\{1,2,3,\ldots\}$ and are told that there is a uniform distribution of probability over the set. For generality, note that $P(\mathbf{N})$ need not equal one (other values may be possibly but we are ignoring them for now). From this given information and nothing more, derive the probability distribution of ODD={1, 3, 5, …). For a given probability distribution over a set, it should be possible to derive the distribution of a (reasonable) subset of elements from the given distribution. A practical understanding of probability should at least give us this. While one can use a symmetry argument to answer this question, such as $P(\mathbf{N})$=P(ODD)+P(EVEN)=2P(ODD) where I am indifferent to ODD and EVEN and thus assign equal probability; I should not have to do this. In finite examples we do not invoke symmetry arguments to derive the probability of a subset of elements – we simply calculate it. What about more difficult infinite examples? For instance, given a uniform distribution over $\mathbf{N}$, what is the probability of $A=\{1, 4, 6, 9, 11, 14, \ldots\}$? It is not clear how symmetry can help us here, at least not trivially, yet intuitively A should have some finite probability of occurring given a uniform distribution over $\mathbf{N}$. Both A and $\mathbf{N}$ have countably infinite members, so we cannot argue that the size of A is less than $\mathbf{N}$.

For finite discrete examples of probability it is common to talk about the number of alternatives corresponding to events, but this is not the case for continuous examples of probability. We understand continuous probability using the theory and language of density functions and integration. The number of alternatives has little or no meaning in continuous probability because most typical intervals over the real numbers have an uncountably infinite number of elements denoted by the continuum or cardinal number $\aleph_1$. Although each interval has uncountably infinite possibilities, it is still possible to assign more or less probability to a particular interval via the density function over that interval. The density function tells us in some sense about the 'spacing' of elements within a particular neighborhood of the real line, even though it says nothing about the number of elements. Density is how we differentiate uncountably infinite sets on the real line, and perhaps we can analogously apply a similar principle to countably infinite sets. For illustrative purposes, in Figure 1 on the left I show a schematic diagram of a continuous probability density function on the real line. The point is that two intervals may each contain an uncountably infinite set of elements but be measured differently because of the density of each interval.

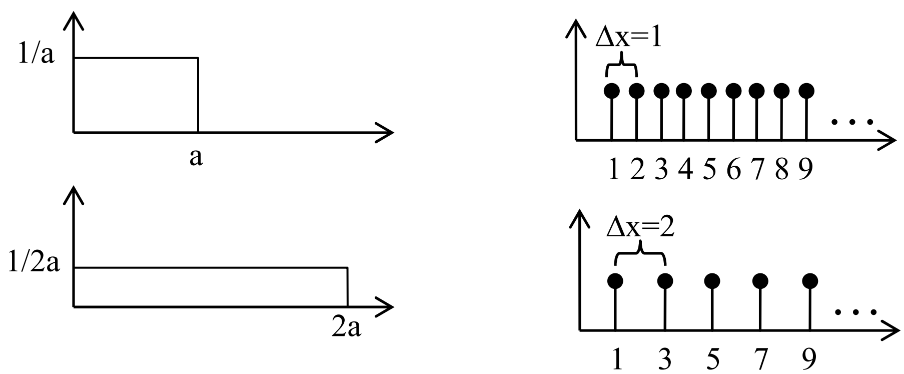

Consider again the task of deriving the probability of ODD from **N**. It is obvious that **N**={1, 2, 3,…} and ODD={1, 3, 5,…} differ in some way, but it is not in the size of the sets. Whereas the sizes are equivalent, we may argue with an analogy to density functions that the relative 'packing' or 'density' of elements differ between the two sets. Looking at the right hand side of Figure 1, it seems reasonable to say that the elements of ODD are half as densely packed as the elements of **N**. We can therefore infer that the probability measure over ODD will be half that of **N**, for ODD contains half of **N**'s density. Again, why do I argue in this way? The answer is that I wish to be able to derive ODD's probability from **N** rather than invoke symmetry arguments that seem unneeded in the finite and uncountably infinite case. Further, the concept of density applied to countably infinite sets has practical utility. Be aware that this notion of density is a relative measure that can only be defined relative to a given sample space of interest.

Example 1. Given a uniform distribution over **N**={1, 2, 3, …}, what is the probability of selecting at random an element x from the subset A={1, 4, 6, 9, 11, 14, …}? We can decompose A into A_1={1, 6, 11, …} and A_2={4, 9, 14, …}. Each of these infinite sets has 1/5 the density of **N**, therefore $P(x \in A) = 2/5$.

Example 2. Given a uniform distribution over **N**={1, 2, 3, …}, what is the probability of selecting at random an element x from the subset A={1, 10, 100…}? The mean interval between the elements in A is infinity, therefore A is infinitely less dense than **N** and $P(x \in A) = 0$.

Example 3. Given a uniform distribution over a countably infinite set B that contains elements separated by a mean interval $E(I_B)= \mu_B$, what is the probability of selecting at random an element x from an infinite subset A that contains elements separated by a mean interval $E(I_A)=\mu_A$? $P(x \in A) = \mu_B / \mu_A$. This attempt at a generalization is undefined when μ_B is infinite and also requires that the set B is ordered like the natural numbers. It fails to hold, for instance, on the integers which have a different order-type, although we can reorder the integers as {0, -1, 1, -2, 2,…} and meaningfully apply the density concept as above.

This next example addresses probability over countably infinite sets but only indirectly deals with the notion of density as above. I include it here in the hope that it will encourage us to think carefully before creating and analyzing additional so-called paradoxes to the Principle of Indifference.

Example 4. The re-labeling paradox (Bartha 2004, attributed to John Norton). Briefly, let ONE={1, 5, 9,…}, TWO={2, 6, 10,…}, THREE={3, 7, 11,…}, and FOUR={4, 8, 12, …}. Now select a natural number x at random. Bartha first reasonably assumes that $P(x \in \text{ONE}) = P(x \in \text{TWO}) = P(x \in \text{THREE}) = P(x \in \text{FOUR}) = 1/4$.[4] The paradox is derived by mapping the original outcome space to a new outcome space in three steps, where the mapping f: **N**→**N**' is represented by the following table:

1→1'	2→4'	3→2'
5→3'	4→8'	7→6'
9→5'	6→12'	11→10'
…	…	…

The odd numbers in **N**' correspond to the set ONE in **N**. We assume that $P(\text{ODD}) = 1/2$, but after re-labeling we may also reason that $P(\text{ODD}) = P(\text{ONE}) = 1/4$ since ODD and ONE represent different names for the same set of elements, creating a contradiction.

I do not agree that this is a contradiction or paradox, rather, I see yet another example that all probabilities are conditional on a given sample space (Burock [2005], Hajek [2003]). This re-labeling paradox is no different than Bertrand's chord paradox so many years ago, both of which arise when we begin with a particular sample space and then

[4] I have greatly distorted and ignored the language of relative betting quotients used in Bartha [2004]. Although useful and interesting, they are unneeded in the present formulation while using the language of this paper.

map this *given* space to a new infinite sample space. We are given in the problem that $P(x \in \{1,5,9,...\}) = 1/4$ and then told to map {1, 5, 9,…}→{1', 3', 5',…}. If you have mapped correctly, your probability measure had better give $P(x \in \{1',3',5',...\}) = 1/4$, for these events correspond to the same outcome. If we had not attempted this mapping from the given sample space, then it would be reasonable to invoke the Principle of Indifference and argue $P(x \in \{1',3',5',...\}) = 1/2$. But these are two completely different situations, one that involves a re-mapping from a given sample space and another that assumes a completely different sample space. In this example we should not be surprised that $P(\{1',3',5',...\} \mid \mathbf{N}')$ and $P(\{1',3',5',...\} \mid \mathbf{N} \to \mathbf{N}')$ have different measures of probability under the assumption of Indifference because the two propositions are fundamentally different entities. By what reasoning must Indifference give the same probability to completely different situations? While the Principle of Indifference may yield identical probabilities in some remapped sample spaces (finite spaces for example), it is not required to do so. There has never yet been as some authors say a 'reckless use of the Principle of Indifference', but only the faulty guidance of intuition in propelling us to compare unlike situations.

References

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

Bartha, P. [2004]: 'Coutable Additivity and the de Finetti Lottery', *Brit. J. Phil. Sci*., 55, 301-321.

Burock, M. [2005]: 'Indifference, Sample Space, and the Wine/Water Paradox', *PhilSci Archive*, Retreived Jan 10, 2006, from http://philsci-archive.pitt.edu/archive/00002487/

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

Hajek, A. [2003]: 'What Conditional Probability Could Not Be', *Synthese*, 137, No. 3, 273-323.

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

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

Vallentyne, P. [2000]: 'Standard Decision Theory Corrected', *Synthese*, 122, 261-290.