

Infinite Decisions and Rationally Negligible Probabilities

NICHOLAS J.J. SMITH

University of Sydney

njjsmith@sydney.edu.au

Abstract

I have argued for a picture of decision theory centred on the principle of Rationally Negligible Probabilities. Isaacs argues against this picture on the grounds that it has an untenable implication. I first examine whether my view really has this implication; this involves a discussion of the legitimacy or otherwise of infinite decisions (decision situations in which an agent must decide upon a choice from infinitely many available options). I then examine whether the implication is really undesirable and conclude that it is not.

1 Introduction

In Smith [2014] I argued for a picture of decision theory centred on the principle:

Rationally negligible probabilities (RNP) For any lottery featuring in any decision problem faced by any agent, there is an $\varepsilon > 0$ such that the agent need not consider outcomes of that lottery of probability less than ε in coming to a fully rational decision.

I introduced the following decision method, which flows naturally from RNP:

Truncation When you are faced with a decision problem d that involves a lottery L , pick a probability ε that is rationally negligible with respect to d , L and yourself—by RNP we know there is at least one such—and then set your value for L to your value for L/ε .

Here, L/ε —the ε -truncation of L —is a particular lottery that is as similar as possible to L while assigning probability zero to outcomes to which L assigns probability less than ε .¹ I also defended the following additional principle:

Weak Consistency If a single decision problem d involves two lotteries L_1 and L_2 over the same set of outcomes, then if one truncates L_1 at ε for purposes of addressing d , one must also truncate L_2 at ε for purposes of addressing d .

I did not argue that this picture of decision theory is free from intuitive cost. I explicitly noted that one of the costs of the picture is denying the principle of *evaluative compositionality* (according to which the value that a rational agent places on a gamble is a function of the values that she places on the possible outcomes of the gamble, together with the probabilities assigned to those outcomes by the gamble). Rather, I argued that the picture of decision theory that we get if we adopt RNP is more unified, tractable and plausible than any alternative picture currently available.

Isaacs [forthcoming] argues against the picture of decision theory based on RNP on the grounds that it has an untenable implication. In particular, the picture deems it rationally permissible to take certain gambles which—Isaacs claims—should not be taken.

¹Here, and elsewhere, see Smith [2014] for full details.

The gambles in question are versions of Penny or Doom. Each is decided by flipping a coin until it lands heads for the first time (and then stopping). There is one version— PD_n —for each $n > 0$: PD_n pays out $-\$(2^{3n-2})$ if the coin is tossed exactly n times; otherwise it pays out one penny. Isaacs claims that decision theory based on RNP deems it rationally permissible for an agent to value each PD_i , for i greater than some number, at one penny—and hence to prefer all such PD_i 's to Status Quo (a lottery that pays out \$0 however the coin lands). Isaacs claims that this is 'deeply wrong' and calls it an 'untenable result'; he claims that 'Such gambles [i.e. all the PD_n 's] should not be taken at all' [Isaacs, forthcoming, 3, 4].

For ease of reference in what follows, let us isolate the following claim:

P It is rationally permissible for an agent to value each PD_i , for i greater than some number, at one penny—and hence to prefer all such PD_i 's to Status Quo.

There are two questions that we need to explore. First, is **P** really a consequence of my view of decision theory? Second, is **P** really a bad result? I address these questions in sections 2 and 3 respectively.

2 Does **P** follow?

Isaacs offers two ways to derive **P** from my views. The first appeals to RNP and Weak Consistency. The second appeals to RNP and dominance (and not Weak Consistency).

The second derivation can be dealt with quickly because it contains two decisive flaws. First, it is dialectically ineffective. I argued that one of the advantages of my view of decision theory is that given Weak Consistency, we can vindicate dominance reasoning. I explicitly did not propose dominance as a free-standing principle to be added to RNP—and I did not claim that we can recover dominance in the absence of Weak Consistency. Second, the derivation is in any case mistaken. Isaacs introduces a lottery Bad News. It, like the PD_n 's, is decided by flipping a coin until it lands heads for the first time (and then stopping); the payoff is $-\$(2^{3n-2})$, where n is the total number of times that the coin is flipped. In a situation in which Bad News and all the PD_n 's are to be decided by the very same sequence of coin tosses, each PD_n dominates Bad News. Truncation allows an agent to assign Bad News a finite negative value v . By dominance it then follows that each PD_n must be assigned a value greater than v . But the expected utilities of the PD_n 's decrease as n increases in such a way that some of the PD_n 's cannot then be valued at their expected value: 'infinitely many versions of Penny or Doom need not be valued according to their expectations' [Isaacs, forthcoming, 3]. So far so good. However, Isaacs continues: 'These versions of Penny or Doom must have bad outcomes which may be ignored, thus these versions of Penny or Doom may be rationally valued at a penny. And once again, a penny is better than nothing.' This does not follow. From the fact that each PD_n must be valued at an amount greater than v it does not follow that *any* PD_n must be valued at one penny—nor indeed given a valuation (be it one penny or something else) greater than that of Status Quo. It could just as well be that the versions of PD_n that are not valued at their expected values are valued at some amount slightly greater than v —an amount less than the value of Status Quo.

Let us turn then to Isaacs's first derivation. As presented by Isaacs, the derivation makes reference to Bad News. However, this is unnecessary. We need only consider a

situation in which we are to choose one out of Status Quo and all the PD_n 's (to be decided by the very same sequence of coin tosses). Consider an arbitrary PD_i . Let us—as we are rationally permitted to do, according to my view—truncate it at toss k , for some k (i.e. we ignore outcomes of probability less than the probability of heads coming up first on toss k). k may be greater than, less than or equal to i : it doesn't matter. The key point is that there will be some PD_j —indeed infinitely many—such that $j > k$. By Weak Consistency, we must truncate PD_j at toss k . We thereby ignore its bad outcome and value it at one penny, hence preferring it to Status Quo.

Note that it is essential to this derivation that we are faced with a single decision problem featuring infinitely many options: *all* the PD_n 's.² If there are only finitely many, then while any of them may be truncated, there is no guarantee that the truncation point need come *before* the bad outcome of *any* of the PD_n 's featuring in the decision problem. Thus my view does not have the following as a consequence:

P' It is rationally permissible for an agent to value each PD_i , for i greater than some number, at one penny—no matter in what context those PD_i 's are encountered.

The question is whether my view has the following as a consequence:

P'' In the context of a single infinite decision d that involves all the PD_n 's and Status Quo, it is rationally permissible for an agent to value each PD_i , for i greater than some number, at one penny—and hence to prefer all such PD_i 's to Status Quo for purposes of making decision d .

This raises the crucial question whether infinite decision situations are licit. If they—or at least some of them—are licit, then a second question also needs to be addressed: in Smith [2014], I was thinking only about standard, finite decisions; if we countenance infinite decisions, the question arises whether we should still accept RNP and Weak Consistency in this new setting. I address these two questions in sections 2.1 and 2.2 respectively.

2.1 Infinite decisions

By an infinite decision, I mean a decision situation in which an agent must decide upon a choice from infinitely many available options. The concept of an infinite decision is therefore distinct from both of the following concepts:

1. An infinite *lottery*. Here a chance process produces an outcome and there are infinitely many possible outcomes that it might produce.³ More technically: part of the definition of any lottery is a probability measure; in the case of an infinite lottery, the sample space over which the probability measure is defined has infinitely many elements.
2. An infinite *sequence* of decisions. Here an agent must make a series of decisions, one after the other—and this series continues infinitely. (None of the individual decisions need involve making a choice from more than finitely many available options.)

²The same point applies to Isaacs's derivation, which goes via Bad News.

³The crucial distinction is therefore between a *decision* and a *chance process*—more on this below.

Of course we can have various combinations of the three. For example, we could consider an infinite sequence of decisions, some of which are infinite decisions; and Isaacs's case of the PD_n 's involves an infinite decision and an infinite lottery (the sequence of coin tosses has infinitely many possible outcomes: heads first on toss 1, 2, 3, ...).

Are infinite decisions licit? Isaacs evidently just assumes that they are: he does not discuss the issue. In fact the issue seems not to have received much discussion in the literature. Here is one comment from Arntzenius et al. [2004, 260]:

one might ban decision situations with infinitely many options, require rational agents to have upper bounds on their utility functions, and ignore cases in which agents divide their credence among infinitely many alternatives. We recognize this craven line of retreat as a potential last resort, but as nothing more. For we are loath to constrain the scope of decision theory with such seemingly ad hoc bans.

Arntzenius et al here lump infinite decisions (i.e. 'decision situations with infinitely many options') in with unbounded utilities and infinite sample spaces—and Nover and Hájek [2004, 246–8], in their defence of decision problems (such as their own Pasadena problem) involving infinite sample spaces and unbounded utilities, run a slippery slope argument along the following lines. It may be that concrete decision situations in everyday life are always finite—but decision theory involves idealisation. We are happy with infinity and unboundedness at all sorts of places in our idealised theories—so we should not balk at them here: there would be a slippery slope from banning infinite sample spaces and unbounded utilities to banning infinity and unboundedness in places where we do not want to ban them. In Smith [2014] I explicitly eschewed restricting decision theory to finite lotteries or to agents with bounded utilities. However, I do not think that a slippery slope strategy would suffice to get us from this point to infinite decision problems—because such a strategy fails in a related case.⁴ Consider the case of someone calculating the value of a function for given argument(s) by following an algorithm or effective procedure (e.g. adding some numbers). It may be that concrete situations of this sort are always finite: the person has a finite memory, a finite amount of time available (bounded by her own lifespan, if by nothing else), a finite available supply of paper and pencils, and so on. In a standard and extremely widely accepted idealisation of this situation—Turing machines—we allow some infinities *but not others*. The Turing machine has an infinite tape and an unbounded running time. The length of the input is unbounded: for example, there is no positive integer that is simply too big for the Turing machine to comprehend (whereas there *is* an upper limit on the numbers that an ordinary human can process). However, it is still required that the input to the Turing machine be *finite*. We idealise away the upper bound on the length of the input (which means that the *tape* on which the input appears is taken to be *infinite*) but not its finiteness (i.e. the input itself is taken to be *finite*). Similarly, we idealise away the upper bound on the computation time of a realistic agent but not its finiteness (i.e. we still suppose that the machine is to halt with its output after some finite amount of time)—and likewise for the machine's programme (unlike with a human agent, there is no upper bound on the number of the machine's internal states

⁴To be clear, I am not attributing such a slippery slope argument (from unbounded utilities and infinite sample spaces, to infinite decision problems) to Nover and Hájek or to Arntzenius et al: I am imagining a kind of argument that someone might extrapolate from their papers.

or instructions—but both numbers must be finite).⁵ Thus the idea that when we idealise away some finiteness we must (or may) idealise away it all—on pain of arbitrariness—is incorrect. Hence a slippery slope strategy would fail to show that infinite decisions are licit, even given that infinite lotteries and unbounded utilities are licit.

Furthermore, I think there is a very good reason to hold that at least some infinite decisions are unacceptable (i.e. even in the idealised setting of decision theory—as opposed to realistic everyday situations). Contrast a decision (as made by an ideally rational agent) with a chance process. It is fine for an (idealised) chance process to have infinitely many possible outcomes: all it has to do is, so to speak, *land on* one of them. In order to do this, it does not first have to survey all the possible outcomes and then decide on one on which to land—it just has to land somewhere. A rational decision, however, is very different. I can hardly make a rational decision amongst a range of options if I do not, at a minimum, consider all the options before making my decision. If I just plump for one I may get lucky but I will not have made a rational decision. But then it looks as though there is a very good reason why infinite decisions are *not* OK even in an idealised setting and even though lotteries with infinitely many outcomes are OK in such a setting: in making a rational decision, an agent—whether realistic or idealised—must at least consider all the available options before deciding on one; if there are infinitely many options, it seems that such consideration will preclude ever coming to a decision.

Note that the talk about ‘consideration’ here does not need to be taken too literally. Taking our lead from discussions in epistemology concerning justification, we can distinguish *internalist* and *externalist* accounts of rationality. On the internalist picture, coming to a rational decision involves giving yourself reasons for choosing a certain option—it involves deciding *responsibly*. On the externalist picture, coming to a rational decision need not involve such internally accessible reasons: you might just embody some process that (say) weighs probabilities and utilities—and the first you come to know of your decision at a conscious level is feeling an impulse to choose a certain option. Still, this will be a rational choice—on the externalist understanding—if the mechanism involved is reliable or well functioning in the right kind of way. On either the internalist or the externalist understanding, the point of the previous paragraph goes through: whether each option need be consciously considered or simply processed or weighed by a mechanism, still a rational decision is very different from a chance process. Rational decision involves, at a minimum, considering or weighing all the options: it is not enough just to plump for one without processing them all.

In fact, one might even consider thinking of rational decision making as a matter of following an effective procedure—and model it using a Turing machine. The machine takes as inputs (representations of) the possible courses of action; it gives as output a single course of action (the rational choice).⁶ If the machine does not even traverse all the

⁵Of course there are other situations in which we consider Turing machines that do not halt on certain inputs; and there are situations in which we consider Turing machines that begin on infinite input strings. This is irrelevant to my present point, which concerns the case of Turing machines *as an idealised model of* an agent calculating the value of a function for given argument(s) by following an effective procedure. In this case the realistic situations being modelled are finite, and in the model we idealise away finiteness in certain places—e.g. the amount of working paper available and the lifespan of the agent—but not in others—e.g. we require the answer to the calculation to be finite (as opposed to an infinite string of symbols) and to be delivered after a finite number of steps of computation.

⁶Of course different decisions might be rational for agents who have different beliefs (probabilities)

inputs—does not consider all the possible courses of action—then it cannot be making an ideally rational decision amongst them. But then there cannot be infinitely many inputs—or we shall never get a decision after a finite amount of processing.

In any case, the idea that it involves following some sort of *procedure* that processes the possible courses of action seems crucial to the idea of rational decision making. Imagine an oracle that simply selects a course of action instantly from any number of possibilities (even infinitely many), without following any procedure. It may be that doing what the oracle recommends always makes you happy, fulfilled and so on. This would make the oracle a great thing to have around—but it wouldn't make it an *ideally rational decision maker*. The oracle isn't something you could aspire to emulate in your own decision making: it isn't doing perfectly what you are fallibly trying to do when making decisions. It may have useful *advice* on *what* to choose—but it cannot teach us any lesson about *how* to make a rational choice. It's a magical substitute for rational decision making—not an ideal decision maker.

We have a reason, then, for being suspicious of infinite decisions that does not also apply to infinite lotteries: a foothold on the slippery slope. Nevertheless it seems that perhaps sometimes we might be able to do something like decide among infinitely many options—without ignoring any of them but also without having to process each of them *one at a time*. One standard way of getting a finite handle on an infinite collection, of getting an infinite genie into a finite bottle—a way of considering *all* the members of an infinite collection without having to consider each one individually or separately—is via the twin techniques of recursive definition and inductive proof. If the infinitely many options can be enumerated in such a way that the step from an option numbered n to an option numbered $n + 1$ is always the same kind of step (or at least one of only finitely many kinds of step)—no matter what n is—then by describing the options numbered 1 and the kind(s) of step required to get from an option numbered n to an option numbered $n + 1$ we describe all the options in a surveyable, manageable way (or at least, in a finite way). This kind of approach is ubiquitous in logic, for example—but we can also see how it would help in decision theory. For example, suppose that a die is to be thrown. Let i be the result of the throw (so i is a number between 1 and 6—but we don't yet know which number). We are offered infinitely many deals and we are to pick one—but the deals are related in an orderly way: there is one deal for each natural number n ; the price of the deal is $f(n)$; the payoff of the deal is $g(n, i)$. (Imagine that the functions f and g are actually described by the person offering the deals.) Now suppose that we can express the expected return of deal n in terms of f and g (and assuming the die is, say, fair). Then we may be able to make this decision by finding a value of n that maximises the expected return. If so, we can thereby make a rational choice among all the infinitely many options available, without ignoring any of the options but also without having to consider them one by one. Compare: solving the equation $x + 5 = 100$ is quite different from picking a natural number randomly—and yet we can do the former without having to consider each possible x (each natural number) one by one (or even each x from 1 through 95).

Of course this sort of approach will only work when the infinitely many options are related in an appropriately orderly way. Suppose we are offered a different infinite set of deals—each involving paying a certain amount of money in exchange for a certain kind of animal—and we are to choose a deal. Suppose that the kinds of animals involved in

and/or desires (utilities). There are various options here—but the details are irrelevant for present purposes.

the first twelve deals are as follows: (i) belonging to the emperor, (ii) embalmed, (iii) tame, (iv) sucking pigs, (v) sirens, (vi) fabulous, (vii) stray dogs, (viii) frenzied, (ix) innumerable, (x) drawn with a very fine camelhair brush, (xi) having just broken the water pitcher, (xii) that from a long way off look like flies.⁷ There is no rule linking one option to the next. To make such a decision in a rational way (as opposed to picking an option at random), we would need to consider each option individually. Thus, if the list of options is infinite, we cannot make the decision in a rational way in a finite amount of time.

Let's take stock. We have seen a reason to be suspicious of infinite decisions that does not also apply to infinite lotteries. However, for all we have said, some infinite decisions are OK. Given some natural assumptions, by a simple counting argument there are *more* illicit infinite decisions than OK ones—but this point need not detain us here, because the infinite decision that Isaacs presents is one of the OK ones. The infinitely many PD_n 's can be arranged in such a way as to allow us to consider them all without considering them one by one. In fact we already did precisely this when we indexed them by $n \in \{1, 2, 3, \dots\}$ and said that PD_n pays out $-\$(2^{3n-2})$ if the coin is tossed exactly n times and one penny otherwise. We have, then, arrived where Isaacs started: at the point of taking a certain infinite decision to be worthy of consideration. Nevertheless we have made genuine progress. For we have now arrived at this point in a responsible way—rather than by sheer assumption or, for example, by a fallacious slippery slope argument.

2.2 Infinite decisions and truncation

Given our conclusion that at least some infinite decisions are licit, we now need to address the question whether we should still accept RNP and Weak Consistency when infinite decisions are involved (for when I advocated these principles in Smith [2014], I was thinking only about standard, finite decisions).

We can derive an argument for rejecting RNP in the context of infinite decisions from Hájek [2014, 564–5]:

When we are giving advice to humans for their everyday decisions, it would be madness to attend to all probabilities, however small. We are boundedly rational, physically impoverished agents constrained by limitations in cognitive ability, time, powers of discrimination, the sensitivity of our experiments, and so on. For us, it is rationally permissible, and even rationally required, to ignore sufficiently small probabilities. . . . However, when we are discussing the St Petersburg game and the Pasadena game, we have already signalled that we are up to our necks in idealization . . . We have left the constraints of the actual world far behind. Our topic is the ideal response to a highly idealized thought experiment about a physically impossible game. If you insist on tolerance here, then I would say that you are changing the topic. You are simply not playing the game.

The argument is as follows. RNP applies only in everyday decisions—and it applies there because of limitations on our time, cognitive abilities and so on. When we are considering

⁷This list is adapted from Foucault [1970, xv], who attributes it to Borges.

infinite decisions, we are clearly considering idealised situations—and so RNP no longer applies.⁸

I do not wish to run such an argument, however—for I did not advocate RNP as a principle that applies only to ordinary decision makers in ordinary decision situations, due to limitations on time, cognitive abilities and so on. Rather, I argued that even though we sometimes *can* pay attention to smaller and smaller probabilities *ad infinitum*—this need not involve performing a supertask and need not be practically impossible—nevertheless, beyond a certain point, factoring in outcomes of lower and lower probability does not make one’s decision any *better*, any more *rational* [Smith, 2014, 473–5]. I argued for RNP in the context of a more general discussion of the need for tolerances in normative theories of practical activities and said that ‘the point is not that arbitrary precision is unattainable: it is that no-one should *require* it. *That* is why a normative theory must not demand infinite precision’ [Smith, 2014, 473]. The tolerances are thus part of the *normative* theory. So the idea was that RNP applies not only to ordinary humans in practical decision situations but also to an ideal decision maker engaged in a practical decision situation or in a legitimate idealisation of such a situation.

With this point in mind, I think there is no reason to reject RNP in the context of infinite decisions. RNP was supposed to apply not only in realistic, everyday decision situations but also in legitimate idealisations of such situations involving ideal decision makers. Given that some infinite decision situations are legitimate idealisations of real decision situations, I see no reason why RNP should not apply in such situations.

What about Weak Consistency? Isaacs [forthcoming, n.9] writes: ‘Weak Consistency implies that in a decision problem involving any number of lotteries (even infinitely many) if an agent truncates one of the lotteries at ϵ that agent must truncate all of the lotteries at ϵ . Repeated applications of [Weak Consistency] easily secure this result.’ Whether or not Isaacs’s claim is true depends on how we read Weak Consistency (and Isaacs’s claim). If we read Weak Consistency as talking about a decision problem that involves exactly two lotteries then Isaacs’s claim is false.⁹ If we read Weak Consistency as talking about a decision problem that involves any finite number of lotteries, of which we consider any two, then again Isaacs’s claim is false. If we read Weak Consistency as talking about a decision problem that involves any number of lotteries—finite or infinite—of which we consider any two, and if we take it that the lotteries involved in Isaacs’s claim are defined over the same set of outcomes, and if we read ‘repeated’ in Isaacs’s claim as ‘infinitely many’, then his claim is true. However, none of this really matters for present purposes. For suppose that we simply replace Weak Consistency with the (possibly) stronger principle that Isaacs needs for his argument (rather than trying to derive it from Weak Consistency):

If a single decision problem d involves any number of lotteries (even infinitely many) over the same set of outcomes, then if one truncates one of them at ϵ for purposes of addressing d , one must also truncate all of the others at ϵ for purposes of addressing d .

⁸In fact Hájek thinks that we enter the idealised realm as soon as we consider the St Petersburg and Pasadena games.

⁹Compare the following case (which should be salient in discussions of the Pasadena game!): from the fact that for any two numbers a and b , $a + b = b + a$, it does not follow that infinitely many numbers yield the same sum no matter the order in which they are added.

Given that we accept certain infinite decisions as legitimate, I think that the reasons presented in Smith [2014] for accepting Weak Consistency also give us reason to accept this principle.

3 Is P an undesirable result?

We have found no reason to reject the claim that my view of decision theory has **P**—or more precisely **P'**—as a consequence. Thus we come to the question: is this really an undesirable result? I shall argue that it is not.

Isaacs glosses the result as licensing ‘taking arbitrarily much risk for arbitrarily little reward’ (this formulation appears four times in his paper and once more in the abstract). Now I think we can probably all agree that taking arbitrarily much risk for arbitrarily little reward would be irrational—but I reject this formulation as a gloss on **P'**. If someone is prepared to take arbitrarily much risk for arbitrarily little reward, that means that no matter how much we lower the probability of winning and/or lower the value of the prize and/or increase the undesirability of losing, she will still be prepared to take the gamble. For example, suppose that Bob will jump off the local bridge into the river for \$20. In fact, he would do it for \$10—or \$5, or \$2.50 . . . He would also jump off the bridge even if it were 1m higher—or 2m, or 3m . . . Bob is then prepared to take arbitrarily much risk for arbitrarily little reward. This is very different from what **P'** allows. **P'** permits ignoring the possibility of the bad outcome of infinitely many PD_n 's (for purposes of making a single decision involving all the PD_n 's). As n increases, the badness of the bad possible outcome of PD_n increases—but its probability decreases. Thus, being prepared to accept any of these PD_n 's over Status Quo is not ‘taking arbitrarily much risk’. In one sense it is: you are risking worse and worse bad outcomes. But in another sense it is the exact opposite: the probability that anything other than a good outcome will occur gets less and less as n increases. As for ‘arbitrarily little reward’, there seems to be no sense in which this description is apt: the reward (should one get it) is constant (one penny) as n increases. So, as n increases, being prepared to accept PD_n over Status Quo means taking on a *higher and higher* chance of getting a *fixed* reward (one penny). It does also mean accepting a risk of worse and worse bad outcomes—but the *probability* of getting a bad outcome *decreases* as n increases. This, then, is *not* an example of what is ordinarily meant by ‘taking arbitrarily much risk for arbitrarily little reward’.

Putting aside the mistaken gloss, then, let's look at what **P'** licenses and consider whether it is really so bad—or indeed bad at all. The core idea is that one is *not* rationally required to factor in *all* the possible outcomes of the sequence of coin tosses—which have successively lower probabilities—*ad infinitum*. Hence there will be some (indeed infinitely many) PD_n that you can rationally regard as offering a definite penny. But n may be *extremely* large. Name the highest number you can; let's call it a .¹⁰ It does not follow from my view that the bad outcome of PD_a may be ignored. It follows from my view that there is a threshold below which probabilities may be treated as zero—but the threshold might be *much* lower than the probability of first getting heads on toss a . Now the idea that it is *possible* to make n large enough so that PD_n *may* rationally be treated as

¹⁰See Aaronson [1999] for some hints on how to name *really* big numbers. Where c is the number of atoms in the visible universe, c^c is peanuts compared to some of the numbers Aaronson discusses.

if it offers a definite penny simply doesn't sound bad to me! However, rather than simply thump the table about this case, let's consider the alternative—for as we shall see, it is extremely counterintuitive. Let us suppose that we change the definition of the PD_n 's in such a way that the positive payoff is not one penny but some much higher amount $\$m$ —enough to solve all your financial problems (and all the financial problems of everyone you care about, and so on). Of course we also make the bad outcome much worse, so that the expected value of the lottery is still negative and indeed the expected values of the PD_n 's still decrease as n increases. Now, the bigger n is, the lower the probability of getting the bad outcome and the higher the probability of getting $\$m$. Someone who thinks \mathbf{P}' is untenable must also think that it is irrational to prefer *any* of these new PD_n 's to Status Quo—no matter how high the probability of winning and no matter how big we make m . However, it seems to me that for many people there would be an n and an m high enough for them to accept one of these new PD_n 's—and (more importantly) that they would be rational to do so. After all, such a PD_n might be much *better* than any free gift in everyday life. Whenever you accept a gift, there is always some (very) small risk of a bad outcome for you (you might get hit by a bus on the way to pick up the gift; you might—to use an example from Hájek [2014, 547]—quantum tunnel to Alpha Centauri; and so on). If we make n big enough, then the risk of a bad outcome when you accept PD_n could be much less than any of the risks you unthinkingly accept in the course of everyday life. Thus we can in effect make PD_n as certain, as risk free, as something that we would normally call a *free gift* of $\$m$ —and we can make m as big as we like! To claim that it would always be irrational to accept any PD_n is, then, highly counterintuitive.

References

- Scott Aaronson. Who can name the bigger number? 1999. URL <http://www.scottaaronson.com/writings/bignumbers.pdf>.
- Frank Arntzenius, Adam Elga, and John Hawthorne. Bayesianism, infinite decisions, and binding. *Mind*, 113(450):251–83, 2004.
- Michel Foucault. *The Order of Things: An Archaeology of the Human Sciences*. Tavistock, London, 1970.
- Alan Hájek. Unexpected expectations. *Mind*, 123(490):533–67, 2014.
- Yoav Isaacs. Probabilities cannot be rationally neglected. *Mind*, forthcoming.
- Harris Nover and Alan Hájek. Vexing expectations. *Mind*, 113(450):237–49, 2004.
- Nicholas J.J. Smith. Is evaluative compositionality a requirement of rationality? *Mind*, 123(490):457–502, 2014.