Vexing Expectations

Harris Nover and Alan Hájek

We introduce a St. Petersburg-like game, which we call the 'Pasadena game', in which we toss a coin until it lands heads for the first time. Your pay-offs grow without bound, and alternate in sign (rewards alternate with penalties). The expectation of the game is a conditionally convergent series. As such, its terms can be rearranged to yield any sum whatsoever, including positive infinity and negative infinity. Thus, we can apparently make the game seem as desirable or undesirable as we want, simply by reordering the pay-off table, yet the game remains unchanged throughout. Formally speaking, the expectation does not exist; but we contend that this presents a serious problem for decision theory, since it goes silent when we want it to speak. We argue that the Pasadena game is more paradoxical than the St. Petersburg game in several respects.

We give a brief review of the relevant mathematics of infinite series. We then consider and rebut a number of replies to our paradox: that there is a privileged ordering to the expectation series; that decision theory should be restricted to finite state spaces; and that it should be restricted to bounded utility functions. We conclude that the paradox remains live.

1. The Pasadena paradox

It's your lucky day. We offer you at no charge the following game, which with winking homage to a more famous game that inspired it, we will call the *Pasadena game*. We toss a fair coin until it lands heads for the first time. We have written on consecutive cards your pay-off for each possible outcome. The cards read as follows:

(Top card)	If the first heads is on toss #1, we pay you \$2.
(2nd top card)	If the first heads is on toss #2, you pay us \$2.
(3rd top card)	If the first heads is on toss #3, we pay you \$8/3.
(4th top card)	If the first heads is on toss #4, you pay us \$4.
⋮	

In general, the *n*th top card informs you that if the coin lands heads for the first time on the *n*th toss, we pay you $\$(-1)^{n-1}2^n/n$, where a negative sign in front of an amount indicates that you pay us that amount.

Rational as you are, you value the game according to its expected utility, the sum of the utilities for each of the possible outcomes weighted by their corresponding probabilities. Identifying (for now) dollar amount with utility, you naively compute the expected utility of the game as an infinite sum. (In section 2 we will review the mathematical background that we presuppose; until then we hope that you either know the relevant mathematics, or will take what we say on faith.)

$$\begin{aligned} EU(\text{game}) &= 2\tfrac{1}{2} - \tfrac{2^2}{2}\tfrac{1}{2^2} + \tfrac{2^3}{3}\tfrac{1}{2^3} + \ldots + (-1)^{n-1}\tfrac{2^n}{n}\tfrac{1}{2^n} + \ldots \\ &= 1 - \tfrac{1}{2} + \tfrac{1}{3} - \tfrac{1}{4} + \ldots \\ &= \ln 2. \end{aligned}$$

The expectation is the familiar alternating harmonic series, whose value is well known. Since $\ln 2 > 0$, you regard the game as favourable, and you agree to play.

It's your unlucky day. By accident, we drop the cards, and after picking them up and stacking them on the table, we find that they have been rearranged. No matter, you say—obviously the game has not changed, since the pay-off schedule remains the same. The game, after all, is correctly and completely specified by the conditionals written on the cards, and we have merely changed the order in which the conditionals are presented. As it happens, the consecutive cards read:

(Top card)	If the first heads is on toss #1, we pay you \$2.
(2nd top card)	If the first heads is on toss #2, you pay us \$2.
(3rd top card)	If the first heads is on toss #4, you pay us \$4.
(4th top card)	If the first heads is on toss #6, you pay us $\$64/6$.
(5th top card)	If the first heads is on toss #8, you pay us \$32.
(6th top card)	If the first heads is on toss #10, you pay us $\$1024/10$.
(7th top card)	If the first heads is on toss #3, we pay you $\$8/3$.
(8th top card)	If the first heads is on toss #12, you pay us $\$4096/12$
⋮	

Now they are ordered so that a single 'we pay you' card is followed by a run of five 'you pay us' cards (corresponding to negative pay-offs for you), this pattern repeated ad infinitum. You now calculate:

$$\begin{aligned}\text{EU(game)} &= 1 + (-\tfrac{1}{2} - \tfrac{1}{4} - \tfrac{1}{6} - \tfrac{1}{8} - \tfrac{1}{10}) \\ &\quad + \tfrac{1}{3} + (-\tfrac{1}{12} - \tfrac{1}{14} - \tfrac{1}{16} - \tfrac{1}{18} - \tfrac{1}{20}) \\ &\quad + \tfrac{1}{5} + (-\tfrac{1}{22} - \ldots\,.\end{aligned}$$

The series has one positive term of the alternating harmonic series followed by five negative terms, ad infinitum. In general, if we rearrange the alternating harmonic series, writing alternately p positive terms followed by q negative terms, the resulting series converges to $\ln 2 + \tfrac{1}{2}\ln(p/q)$. (See Apostol 1967, p. 416.) With $p = 1$ and $q = 5$, this equals $\ln 2 + \tfrac{1}{2}\ln(\tfrac{1}{5}) \approx -0.11$. Thus, the expected utility is apparently negative. The game suddenly looks unfavourable to you.

It's the luckiest day of your life. A gust of wind blows the cards off the table. You pick them up and stack them up again, this time in another order. Obviously, we have no complaint against you, for the cards themselves and thus the game itself again have not changed. But now they read, in order:

(Top card)	If the first heads is on toss #1, we pay you \$2.
(2nd top card)	If the first heads is on toss #3, we pay you $\$8/3$.
(3rd top card)	If the first heads is on toss #5, we pay you $\$32/5$.
⋮	
(12th top card)	If the first heads is on toss #23, we pay you $\$2^{23}/23$.
(13th top card)	If the first heads is on toss #2, you pay us \$2.
(14th top card)	If the first heads is on toss #25, we pay you $\$2^{25}/25$.
⋮	

There are ever-lengthening runs of 'we pay you' cards, interspersed with single 'you pay us' cards. Indeed, the runs of positive pay-offs are always just long enough that your expectation apparently has a very interesting property:

$$\text{EU(game)} = (1) + (\tfrac{1}{3} + \tfrac{1}{5} + \ldots + \tfrac{1}{23} - \tfrac{1}{2}) + (\tfrac{1}{25} + \ldots - \tfrac{1}{4}) + \ldots$$

Each bracket contains a quantity that is at least as large as 1, and this pattern persists forever. So the sum is at least as large as 1 + 1 + 1 + … . That is, the expected utility is infinite. To your utter delight, you think that the game is not only favourable to you; it is better than any finite reward, however large. It is as good as it gets.

It is the unluckiest day of your life. Clumsy fellows that we are, we drop the cards again. Reassembling them this time, they form ever-lengthening runs of negative pay-offs. And worse than the second ordering, this time the resulting expectation series can be bracketed so that each bracket contains a quantity smaller than −1, with this pattern persisting forever. So the sum is at least as small as −1 −1 −1 … . That is, the expected utility is seemingly *negative* infinity. The game is worse than any finite punishment, however large. It is as bad as it gets.

And so it goes. In fact, it seems that any expected utility whatsoever can be realized by some rearrangement of the cards—more on that shortly. What, then, is the value of the game? We already knew that expectations can misbehave in games with infinite state spaces and unbounded pay-offs such as ours—Bernoulli taught us that long ago with his St. Petersburg game.[1] But vexing though that game is, the Pasadena game is *worse* in several ways.

While we may disagree with decision theory's apparent verdict about how good the St. Petersburg game is—we may not think that it is *infinitely* good—still we should all agree that it *is* good. Certainly it is worth playing if there is no entry fee, for example. And however good it is, this is fixed once and for all, independently of how the game is presented. By contrast, there is no fact of the matter of whether the Pasadena game is good or bad, and indeed we can make it appear

[1] In the St. Petersburg game, we toss a fair coin until it lands heads for the first time. The longer it takes, the better for you. You receive exponentially escalating pay-offs according to the following schedule:

First heads on toss	**Probability**	**Pay-off**
1	$\frac{1}{2}$	\$2
2	$\frac{1}{4}$	\$4
3	$\frac{1}{8}$	\$8
⋮	⋮	⋮
n	$\frac{1}{2^n}$	$\$2^n$
⋮	⋮	⋮

Your expectation (in dollars) $= 2\tfrac{1}{2} + 4\tfrac{1}{4} + 8\tfrac{1}{8} + \ldots$
$= 1 + 1 + 1 + \ldots$
$= \infty$

Decision theory is apparently telling you that you should pay any finite amount to play this game once. This seems absurd—and thus we have the St. Petersburg paradox.

exactly as good or as bad as we like by merely presenting the pay-offs in a suitable order. Yet since it is the very same game that is being presented one way or another, it seems that the game is simultaneously good, bad, neutral, incredibly wonderful, unspeakably awful etc.—in short, all things at once. We may put this paradox geometrically, since the expectation of a game corresponds to the point of balance of a 'see-saw' that has weights equal to the probabilities sitting at positions on the real line corresponding to the pay-offs. The see-saw whose positioning of weights corresponds to the St. Petersburg game has no finite point of balance, 'tipping to the right' wherever we put its fulcrum. But the see-saw corresponding to the Pasadena game can apparently be made to balance anywhere by suitable rearrangement of the pay-off table.

The Pasadena game also seems to allow us to set up the nastiest of money pumps: sell the game at a high price, and buy *the very same game* at a low price, with the prices dictated by the putative corresponding expected utilities. Dutch Books exploit inconsistent valuations by an agent of one and the same state of affairs, putatively revealing a defect in that agent's state of mind. But there is nothing defective in the above mathematics: the inconsistent valuations are not the fault of any agent, but rather of the game itself coupled with the usual expectation formula. At least with the St. Petersburg game we knew what we were getting into; but the Pasadena game is a supreme Aladdin's genie among games, capable of transmuting itself into whatever one wants it to be.

Finally, in response to the St. Petersburg paradox, there is something to be said for the bullet-biting response: 'the game *should* be valued infinitely, and any intuition to the contrary should be dismissed as an artifact of our finite minds not fully appreciating the true nature of the game; we should learn to live with decision theory's verdict'. (Cf. Martin 2001, Clark 2002.) But no such reply is available in the Pasadena paradox. There is no verdict that we should learn to live with, because decision theory delivers *so many* verdicts—which is to say that it goes silent.

It is an uncomfortable silence. For intuition tells us—indeed, yells at us—that we can make meaningful comparisons between the Pasadena game and other games. It is clearly worse than the St. Petersburg game, for starters. It is clearly worse than a neighbouring variant of the game—call it the *Altadena game*—in which every pay-off is raised by a dollar. (Notice that the Altadena game has all the problems of the Pasadena game.) And the Pasadena game is clearly better than a 'negative'

St. Petersburg game, in which all the pay-offs of the St. Petersburg game are switched in sign. Yet expected utility theory can say none of this.

Ironically, while the expectation in the St. Petersburg game at least gave us a warning that all might not be well—the series diverges, after all—the expectation in the original presentation of the Pasadena game apparently converged to an innocent ln 2 (that is, about 0.69). And yet we could at least say *something* about the value of the St. Petersburg game—we have already said that it is positive—whereas once we see the differences that rearrangements can make, we realize that we can say absolutely nothing about the value of the Pasadena game. The game is apparently well defined, and yet decision theory cannot handle it. Something has to give—either the game itself, or decision theory.

2. Infinite series: review and discussion

Expected utilities are straightforward when there are only finitely many states. So let us confine our attention to decision problems with a countably infinite number of states. Expected utilities are then infinite series. It will thus be useful for us to review briefly some definitions and results from real analysis regarding infinite series.

Consider an infinite sequence of real numbers $a_1, a_2, a_3, \ldots, a_n, \ldots$. We say that this sequence has the *limit L* if, for every $\epsilon > 0$, there exists N such that for all $n > N$, $|a_n - L| < \epsilon$. In less technical terms, we mean that for any tolerance, there exists a point in the sequence such that every element past that point is within that tolerance of the limit. If the limit condition is met, then we write that $\lim_{n\to\infty} a_n = L$. If a sequence has a limit, we say that the sequence is *convergent*. Otherwise, we say that it is *divergent*.

Now, instead of the sequence above, consider a related sequence : We

$$a_1, a_1 + a_2, a_1 + a_2 + a_3, \ldots, \sum_{i=1}^{n} a_i, \ldots .$$

call such a sequence that is made up of the partial sums of the terms of an infinite sequence an *(infinite) series*. Such a series may of course have a limit, by which we mean a number S such that $\lim_{n\to\infty} \sum_{i=1}^{n} a_i = S$. For simplicity's sake, we will refer to both a series and its limit (if one exists) by $\sum_{i=1}^{\infty} a_i$. We classify series just as we do sequences: a series may be *convergent* or *divergent*. The harmonic series $\sum_{i=1}^{\infty} 1/i$ is divergent; but the alternating harmonic series is convergent (and has limit ln 2, as we have noted).

The question of divergence or convergence is decided on the basis of the behaviour of the sequence of the partial sums. It is important to note that the limit of this sequence does not necessarily behave as we would expect from our experience with finite sums. Any reordering of a finite sum yields the same result. However, as we saw in the expectation calculations of the Pasadena game, reordering the terms in an infinite sum may change its value.

There is a condition on convergent series that is both necessary and sufficient for such mischief being possible. We say that the infinite series $\sum_{i=1}^{\infty} a_i$ is *absolutely convergent* if the related infinite series $\sum_{i=1}^{\infty} |a_i|$ is convergent. In this case, $\sum_{i=1}^{\infty} a_i$ will also converge, and the value to which it converges is independent of the order of the terms. Alternatively, if $\sum_{i=1}^{\infty} a_i$ converges but $\sum_{i=1}^{\infty} |a_i|$ diverges, we say that the series is *conditionally convergent*. We can now state the necessary and sufficient condition: rearrangement can change the value of a convergent series iff the series converges conditionally (see Apostol 1967, p. 412).

That does not yet settle how *many* different values can be achieved by rearrangement. In fact, there are uncountably many such values; indeed, a conditionally convergent series can be rearranged to give whatever sum you want. This is Riemann's celebrated rearrangement theorem: *Let* $\sum_{i=1}^{\infty} a_i$ *be a conditionally convergent series of real terms, and let* S *be a given real number. Then there is a rearrangement of* $\sum_{i=1}^{\infty} a_i$ *that converges to the sum* S. (For proof, see Apostol 1967, p. 414.) Moreover, the series may also be rearranged to diverge to positive or negative infinity; and it may be rearranged so that the partial sums forever oscillate without approaching a limit.

3. Infinite decision problems

In a decision problem, the expectation of a given action is a sum. When the set of states that receive positive probability is infinite, it is an infinite sum. So the theory of infinite series applies, and the results above are relevant. However, there is a twist. In real analysis, the order of a particular series is given—there is a fact of the matter of the arrangement of its terms, and any rearrangement of a particular series produces a different series. But in decision theory, there is an arbitrariness to how the pay-off schedule of a given decision problem is presented, and any rearrangement of the schedule yields the same decision problem. (We noted that reordering the cards that detail the Pasadena game does not change the game.) If we let the arrangement of the schedule dictate the arrangement of the corresponding expectation series, we then should regard the choice of the latter to be equally arbitrary. Unfortunately, as we have seen, this arbitrary choice can have serious consequences for the apparent choice-worthiness of a given action. Formally speaking, an expectation that is not absolutely convergent is undefined. But it is all very well to speak formally. Decision theory is a theory of rational action; what is one supposed to *do* in such a case? Should one play the Pasadena game or not, if given the choice? Should one take the Altadena game rather than the Pasadena game, if given the choice? Decision theory, as we have said, is silent.

Consider an action whose expected utility is conditionally convergent. It would already be troubling if there were just two possible values that could be achieved by rearrangement of the expectation series of some action—then, there would be no fact of the matter of the value of that action, a kind of indeterminacy. Still, we might at least be able to make useful comparisons—e.g., if both values were below that of another action, the latter action should be preferred. But the trouble is far worse than that: by Riemann's rearrangement theorem, every possible value (including the 'values' of positive and negative infinity) can be attained by rearrangement. This is indeterminacy on the grandest scale.[2]

One might argue that there is in the end a fact of the matter as to the ordering of the terms: the game is played in a certain sequence, and the ordering thus induced is therefore privileged. We have several misgivings about this argument.

First, it departs from the standard decision-theoretic framework. To determine the choice-worthiness of an action, we should only need to know the probabilities and pay-offs associated with each state of the world under that action. Extra details about how they arise are extraneous to the decision problem. Now, apparently, we are requiring additional information. Yet if other information were to affect choice-

[2] Nathan (1984) is aware that there are games whose expectations may depend on the ordering of terms (although he does not give an example of such a game, and he does not appear to regard them as paradoxical). Nevertheless, he believes that such a game may have a well-defined expectation. He writes:

> In infinite games, we can always proceed by first defining a finite game truncated at step $k = N$ and subsequently letting $N \to \infty$. A complete description of such a game must, therefore, include a method of truncation unless the game has an expectation which is independent of the ordering of the terms of (1) [the expectation series] … If expectation series (1) neither converges absolutely nor tends absolutely to plus or minus infinity, the manner of (virtual) truncation must be specified to meet [the condition that the expectation of the game is uniquely defined]. If this fixes the value of (1), then (1) is the expectation … (pp. 131–2)

In the Pasadena game we could define a sequence of truncated games in which we call the game off if the coin is tossed n times, and then let $n \to \infty$. The sequence of expectations of these finite games has limit ln 2. Does this mean that the desirability of the Pasadena game is, or at least can be, defined after all? We think not. Fixing a method of truncation seems to be no different from fixing an ordering, and we next demonstrate why this cannot be done in a satisfactory manner.

worthiness, it should really be incorporated into either the probability or pay-off profiles.

Furthermore, we chose to use a coin in our game for dramatic effect, but we could have left obscure the details of the mechanism. We could just as easily play with a 'black box'. All you know about the box is that, for every $n \geq 1$, there is a $1/2^n$ chance of the number n occurring, and if that number occurs then you receive a pay-off of $(-1)^{n-1}2^n/n$. Is there still a privileged ordering? That is not at all clear. While n seems to serve naturally as an index here, and thus to provide an ordering, we could just as easily carry out some mathematical transformation on n that would reorder the terms, and present the game that way. For example, we could offer you a game where we index triplets of outcomes. For every $m \geq 1$, there is one pay-off of $\$\frac{2^{4m-3}}{4m-3}$ that has a probability of $\frac{1}{2^{4m-3}}$, one pay-off of $\$\frac{2^{4m-1}}{4m-1}$ that has a probability of $\frac{1}{2^{4m-1}}$, and one further pay-off of $\$-\frac{2^{2m}}{2m}$ that has a probability of $\frac{1}{2^{2m}}$. This is a fairly transparent description of the Pasadena game, with $4m-3$, $4m-1$, and $2m$ replacing n. However, by indexing by m, we are presenting the Pasadena game in a new order, which corresponds in the 'n-notation' to an order of 1,3,2,5,7,4,9,11,6,.... Computing the expected value by this ordering yields a value of $3/2 \ln 2$. Yet in this notation this ordering is 'natural' (whatever that means). Note that we did not change the box itself—we simply changed the notation by which we described its outputs. Without knowing the mechanism, there is no reason to say that this description is less 'natural', as perhaps the mechanism implementing the payoffs actually does make this the most obvious ordering. We should be suspicious, moreover, of an appeal to a notion as ill-understood as 'naturalness'; and even if it were well understood, it is dubious that every decision problem can be given exactly one 'most natural' characterization. This certainly goes beyond anything we were taught in our decision theory textbooks! And of course we always could implement a mechanism that makes the m ordering the most 'natural'. In that case, the suggested line of argument leads to one valuing two games with identical pay-off/probability schedules differently, which is clearly absurd.

4. Apologia

The Pasadena game creates a paradox for decision theory. One way to get rid of the paradox would be to *get rid of the game*. It would be comforting to show that the paradox does not arise in the first place, because the problem is simply ill-posed, the Pasadena game incoherent.

There are two possible lines of attack—neither satisfactory, in our opinion.

4.1. Restrict decision theory to finite state spaces

The first response is to balk at the game's infinite state space: the assumption that the coin could land heads for the first time on the first toss, or the second, or the third, *ad infinitum*. More generally, the response is that decision theory should be confined to actions that have finitely many consequences—that is, to decision problems that have finitely many states. What could motivate this response?

On the one hand, it might be purely theoretical considerations: it might be claimed that *as a matter of conceptual necessity*, all decision problems have finitely many states. But then the response strikes us as high-handed: for we have no trouble countenancing infinitely many states elsewhere in our theorizing—in physics, for example. What the response is *not* is *even*-handed: for decision theory is kindred to probability theory, and yet infinite probability spaces are commonplace. Indeed, Kolmogorov's invocation of countable additivity builds infinitude into the very foundations of probability theory. More than that, the usual justification for maximizing expected utility appeals to the laws of large numbers, limit theorems that are premised on the existence of infinitely many trials, and thus infinitely many states (e.g., states of the form ' … occurs on trial *n*', for infinitely many *n*). In short, the response is at odds with the conceptual underpinnings of decision theory itself, which take infinitude of states in their stride, and indeed which could not be reproduced without such infinitude.

On the other hand, the response might be motivated by empirical considerations about human agents: it might be claimed that *as a matter of contingent fact*, all of our decision problems have finitely many states—we simply do not live long enough, or cannot perform experiments fast enough, or do not have adequate powers of discrimination to generate an infinite state space. But even supposing it to be a fact, so what? Such contingent facts are irrelevant for two reasons. First, it is not the facts of the world that matter, but rather what an agent *believes* those facts to be. While it may be true, for example, that a coin cannot be tossed fast enough to determine every possible outcome of the Pasadena (or St. Petersburg) game in a finite amount of time, if anyone believes that this task can be done, then the infinite state space and hence the problem for decision theory remains. (You might believe, for example, that there is a supreme coin tosser who can complete each toss in half the time of the previous toss.) Second, even if no one actually does believe that this task can be done, the problem is still there. Decision theory is not merely about humans, but is also about—perhaps even especially about—ideal rational agents, and it does not follow that our idealized theory of ideal rationality should be so constrained. Again, the response is not even-handed, for we happily idealize away our all-too-human nature elsewhere in our theorizing—in economics, for example. And again, even decision theory and probability theory themselves ignore our finiteness in other respects, assuming as they do that we are logically omniscient, that our preference orderings are infinitely rich, and so on.

Whichever way this response is run—the 'conceptual necessity' way or the 'contingent fact' way—it cannot stably stop at merely banishing infinite state spaces. Arguments parallel to those adduced in favour of finite state spaces would ultimately have us replace decision theory and probability theory with finitist revisions. Indeed, it would hardly be playing fair to stop there, either; we would need a principled reason why the rest of mathematics should not go the same way. But that way lies finitism and its attendant problems. We do not have space here to enter the debate about finitism, except to note that it commits us to a rather radical revision of mathematics. As a resolution of the Pasadena paradox it is surely overkill.

4.2. Restrict decision theory to bounded utility functions

The second line of attack targets the game's unbounded pay-offs. It has famously been used against the St. Petersburg game by Richard Jeffrey (1983). He writes: 'Put briefly and crudely our rebuttal of the St. Petersburg paradox consists in the remark that anyone who offers to let the agent play the St. Petersburg game is a liar, for he is pretending to have an indefinitely large bank' (p. 154). Jeffrey is avowedly putting the objection quickly here, but it is surely a natural objection, pithily stated. He would presumably rebut the Pasadena paradox in a similar fashion. We think, however, that the paradox won't go away so easily.

If Jeffrey's rebuttal turned solely on a consideration of how much money there is in the world, we could quickly sidestep it by rewriting the pay-offs of the Pasadena game in terms of utiles, units of utility abstracted away from the details of how they are realized. To repeat a point made in the previous section: all that matters to decision theory is the utility/probability profile. If the way in which the rewards are realized mattered, then that should already have been taken into account in either the utilities or the probabilities. So Jeffrey's rebuttal really becomes: *all utility functions are bounded*. Paralleling our discussion in the previous section, this claim can be taken in two ways.

On the one hand, it might be claimed that *as a matter of conceptual necessity*, all utility functions are bounded. But why should we believe this? Unbounded functions abound! In particular, they abound elsewhere in our theorizing. Relativity theory does not balk at unbounded space-time curvature; measure theory does not balk at unbounded lengths. Why should decision theory be special in this regard? More tellingly, probability theory comfortably accommodates unbounded functions—unbounded random variables, for example. The imposition of bounds specifically on utility functions thus seems ad hoc.

On the other hand, it might be claimed that *as a matter of contingent fact*, all humans have bounded utility functions. We think that this is not clearly a fact at all, for even a human might have preferences over infinitely many outcomes, specifiable by a finite rule, and for any pair of outcomes be prepared to pay a fixed finite amount to swap the dispreferred for the preferred outcome. For example, you might be prepared to pay a dollar (or utile) to swap m days in heaven for n whenever n is greater than m; if so, you have an unbounded utility function. But even supposing that all human utility functions happen to be bounded, so what? Again, it does not follow that our idealized theory of ideal rationality should be so constrained. After all, we do not eschew *continuous* utility functions, even though as a matter of contingent fact, all humans have finite thresholds for their perception of reward. Much as we impose no bounds on the sensitivity of our discriminations of utility, we should impose no bounds on the utilities themselves. If we are prepared to idealize 'in the small', we should equally be prepared to idealize 'in the large', fully aware that in doing so we pay no heed to our contingent limitations, in both directions. And even if we do acknowledge the limitations of human beings, decision theory is not specifically about human beings. It is about rational decision makers in the abstract. And so we are brought back to the conceptual argument that we dismissed above.

Thus, we submit that the Pasadena game prevails, and that it is a headache for decision theory. If it is going to be cured, some other kind of medicine will be required.[3]

[3] We thank John Broome, Antony Eagle, Terrence Fine, Matthias Hild, Bradley Monton, Andrew Reisner, Roy Sorensen, Manuel Vargas, Peter Vranas, and especially Mark Colyvan, Adam Elga, Branden Fitelson, Ned Hall, Chris Hitchcock, and Daniel Nolan for very helpful comments.

Michelson Science Center 669
Caltech
Pasadena, CA 91126, USA
nover@its.caltech.edu

HARRIS NOVER

Humanities and Social Sciences Division
MC 228–77
Caltech
Pasadena, CA 91125, USA
ahajek@hss.caltech.edu

ALAN HÁJEK

References

Apostol, Tom M. 1967: *Calculus*, Vol. I. New York: John Wiley and Sons, 2nd edn.

Clark, Michael 2002: *Paradoxes From A to Z*. London: Routledge.

Jeffrey, Richard 1983: *The Logic of Decision*. Chicago: University of Chicago Press, 2nd edn.

Martin, Robert 2001: 'The St. Petersburg Paradox'. *The Stanford Encyclopedia of Philosophy* (Fall 2001 Edition). Edward N. Zalta (ed.), URL = http://plato.stanford.edu/archives/fall2001/entries/paradox-stpetersburg/.

Nathan, Amos 1984: 'False Expectations'. *Philosophy of Science*, 51, pp. 128–36.