

Putting Expectations in Order

Alan Baker[†]

In their paper, “Vexing Expectations,” Nover and Hájek (2004) present an allegedly paradoxical betting scenario which they call the *Pasadena Game* (PG). They argue that the silence of standard decision theory concerning the value of playing PG poses a serious problem. This paper provides a threefold response. First, I argue that the real problem is not that decision theory is “silent” concerning PG, but that it delivers multiple conflicting verdicts. Second, I offer a diagnosis of the problem based on the insight that standard decision theory is, rightly, sensitive to order. Third, I describe a new betting scenario—the Alternating St. Petersburg Game—which is genuinely paradoxical. Standard decision theory is silent on the value of playing this game even if restrictions are placed on the order in which the various alternative payoffs are summed.

1. Introduction: The Pasadena Paradox. In their paper, “Vexing Expectations,” Nover and Hájek (2004) present an allegedly paradoxical betting scenario which they call the *Pasadena Game* (henceforth, PG). A coin is tossed repeatedly until it comes up heads. There are alternating positive and negative payoffs, increasing in size the longer the run of tosses before the first heads appears. Imagine writing the payoff schedule for the game on a set of cards, 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 \$4.¹
- . . .

Calculating the value of playing this game by summing the products of the probability and utility of each outcome listed on successive cards, this yields the infinite harmonic series $1 - 1/2 + 1/3 - 1/4 + 1/5 - \dots$. This series converges to $\ln 2$ (i.e., ≈ 0.693). Is this the value of playing PG?

[†]To contact the author, please write to: Department of Philosophy, Swarthmore College, Swarthmore, PA 19081; e-mail: abaker1@swarthmore.edu.

1. The general rule for calculating the pay-off if the first heads lands on the n th toss is $\$(-1)^{n-1}2^n/n$.

Philosophy of Science, 74 (December 2007) pp. 692–700. 0031-8248/2007/7405-0012\$10.00
Copyright 2007 by the Philosophy of Science Association. All rights reserved.

To generate their paradox, Nover and Hájek run through several alternative scenarios in which the cards describing the schedule are accidentally rearranged. Shuffling the deck does not change the game which the cards describe, yet summing the expected utilities in different orders can produce an overall value of any arbitrary amount, including $+\infty$ and $-\infty$. Nover and Hájek conclude that decision theory must therefore “remain silent” on the value of playing PG. They find this paradoxical since—even if a value for the game cannot be fixed precisely—there do seem to be various comparative judgments we can make between PG and other games. (For example, the ‘Altadena Game’—where the payoffs in the above game are each raised by \$1—is surely preferable to basic PG).

2. Paradox Lost? It should be noted that if Nover and Hájek are correct, then the silence of decision theory concerning the value of PG is the least of our worries. Talk of “silence” suggests that the problem is akin to *incompleteness*: decision theory is too weak because it fails to give any guidance about the value of certain well-defined games. In fact, however, the predicament they describe here is more like *inconsistency*. Decision theory does deliver a verdict on the value of PG; the problem is that it delivers multiple conflicting verdicts. Indeed, as Nover and Hájek themselves point out, the conflicting verdicts allow for the setting up of “the nastiest of money pumps: sell the game at a high price, and buy the very same game at a low price” (2004, 241). Inconsistency of this sort is much harder to ignore—both practically and theoretically. And it makes it imperative to determine if and how the argument to this paradoxical conclusion can be blocked.

The bedrock of Nover and Hájek’s argument for the paradoxical nature of PG is the card reshuffling scenario. Their core argument is as follows:

- (1) The infinite deck of cards, D , correctly describes PG.
- (2) Any reshuffling of D (since the cards are all indexed), also produces a correct description of PG.
- (3) Hence, the sum of the result of each reshuffling of D (taken card by card), has equal claim to be the value of PG.
- (4) Hence, decision theory assigns multiple conflicting values to PG.

It is a good rule of thumb (in philosophy and in life!) to beware of people who try to prove anything using a deck of cards. Nover and Hájek have indeed performed a sleight of hand in the above series of moves. Before diagnosing the flaw, let us examine a structurally similar argument whose obvious invalidity ought to give grounds for suspicion concerning the above argument.

Let S be the infinite harmonic series $1, -1/2, 1/3, -1/4, 1/5, -1/6, \dots$

. The following infinite deck of cards, D^* , provides a correct and complete description of S :

- (Top card) 1 at place 1 in the series.
- (2nd top card) $-1/2$ at place 2 in the series.
- (3rd top card) $1/3$ at place 3 in the series.
- (4th top card) $-1/4$ at place 4 in the series.

Summing these cards in the above order yields an infinite sequence of partial sums which converges to $\ln 2$. Consider the following argument:

- (1*) The infinite deck of cards, D^* , correctly describes S .
- (2*) Any reshuffling of D^* (since the cards are all indexed), also produces a correct description of S .
- (3*) Hence, the sum of the result of each reshuffling of D^* (taken card by card), has equal claim to be the sum of S .
- (4*) Hence, mathematics assigns multiple conflicting sums to S .

Clearly the conclusion of this argument is false: the infinite harmonic series, S , has a unique and well-defined sum, namely $\ln 2$. So where does the argument go wrong?

The problem here stems from a failure to distinguish clearly between the series S itself and the *description* of S . Since the cards are indexed (each card being linked to a particular term in S), shuffling their order still allows S to be reconstructed from the description which the reshuffled cards provide. So if all we are concerned with is having enough information to specify S , all reshufflings of the cards are equally 'correct'. Does this imply that the sums of the card values under these various reshufflings are on a par? No, because shuffling the order of description amounts to rearranging the order in which we add together the individual terms of the series. (It is as if we summed S by taking the first term, then the second, then the fourth, then the third, then the sixth, then the eighth, etc. This is effectively summing a *different* series). For infinite series such as S , a correct description in a gerrymandered order may not yield a correct sum for the series. Why not? Because *order matters* in summing infinite series.

Returning to Nover and Hájek's reshuffling argument for PG, we can now see precisely how it falls short. If we presuppose that order does not matter in assessing the value of games then the reshuffling argument goes through, but it cannot be used to support the conclusion that order does not matter. To derive the 'multiple conflicting values' conclusion of their paradox, Nover and Hájek need an independent argument for why the ordering of games such as PG is not a factor in the decision-theoretic calculation of their overall value.

3. The Importance of Order. Let me put my own cards on the table (no pun intended). I want to defend the claim that order has a role to play in decision theory, and to propose the following principle as an independently plausible way to dissolve the Pasadena Paradox:

Principle O. The expected value of an option that has a countable infinity of different outcomes is equal to the products of the probabilities and payoffs of these outcomes *summed in the same order in which the outcomes occur.*

If we adopt the above principle then we must use the original ordering of the deck (where the top card refers to the first coin toss, the second card to the second toss, etc.) in calculating the expected value of PG.² Thus Principle O ensures that decision theory delivers a unique verdict ($\ln 2$) on the value of playing PG. Nover and Hájek consider a response along the above lines, noting that “one might argue that . . . the game is played in a certain sequence, and the ordering thus induced is therefore privileged” (2004, 244). However, they reject this approach as a solution to their paradox, citing two principal worries.

Nover and Hájek’s first objection is that the formalism of decision theory is not sensitive to order. They argue that taking into account the ordering of the game when calculating the expected value of PG

departs from the standard decision-theoretic framework. To determine the choice-worthiness of an action, we should only need to know the probabilities and payoffs associated with each state of the world under that action. (2004, 244)

Their argument can be broken down as follows: Within the framework of standard decision theory,

- (1) The choice-worthiness of an action is determined by its expected utility.
- (2) Expected utility is equal to the sum of the products of probabilities and payoffs for each state of the world.

2. Principle O applies most straightforwardly to situations in which the various outcomes are independent of one another. Some care is needed in defining a suitable notion of ‘order’ for a game such as PG in which the outcomes are not independent. In what sense does ‘the first heads is on toss #2’ come *earlier* than ‘the first heads is on toss #3’? One way to generate an ordering here is to count the number of constituent choice-events (2 coin tosses versus 3 coin tosses). Or, more complicatedly, one could stipulate that outcome A is earlier than outcome B if and only if before the final coin toss which determines A, the possibility of B is still open, but not vice versa.

- (3) Hence, choice-worthiness depends only on individual probabilities and payoffs.

However, this argument is only valid if the following supplementary premise is added:

- (2*) A sum of terms depends only on the value of each individual term.

The problem is that (2*) simply does not hold in general. It fails for conditionally convergent series, whose sums depend not only on the value of each individual term in the series but also on the order in which the terms appear. Thus it is incorrect to assert that decision theory pays no attention to order. Decision theory utilizes the mathematical operation of summation, and the concept of sum (including the Σ notation) presupposes that the items being summed are ordered in some way.

Nover and Hájek's second objection is epistemological. If the actual ordering of a game implementation cannot be determined by the player, then she cannot apply Principle O. They consider a scenario where the mechanism underlying PG is hidden inside a 'black box'. All you know as a player are the payoffs and probabilities: for each $n \geq 1$, there is a $1/2^n$ chance of winning $(-1)^{n-1}2^n/n$ dollars. You have no idea what is actually going on inside the box. Nover and Hájek ask, "Is there still a privileged ordering?" (2004, 245). The response, surely, is that there is, even if it remains unknown to the player. If the game is ordered then it is ordered. Placing the mechanism which implements the game inside a black box may mask the actual order from the player. This epistemic indeterminacy may make it impossible for the player to apply Principle O, but it does not threaten the cogency of the Principle itself.

Perhaps Nover and Hájek's real worry here concerns not the cogency of Principle O but its plausibility. If Principle O is correct then it is impossible for a player to determine the value of PG without knowing something about the mechanism inside the black box. But—they might argue—how can *that* matter? Once you know the payoffs and probabilities, surely you know all you need to know that is relevant to the value of the game. But once it is phrased this way, it seems clear that the black box argument is not really an argument at all, but simply a restatement of their core claim that order *per se* makes no difference to decision theory.

Interestingly, the black box scenario does—I think—contain the seeds of a more serious objection against Principle O, although this line of attack is not pursued by Nover and Hájek. What if the game inside the black box has *no* intrinsic ordering at all? In other words, what if the indeterminacy of the game mechanism ordering is not merely epistemological but metaphysical? One possibility is to impose an order on the game based on the most natural way of indexing it. In the above example,

n serves naturally as an index, so why not sum the alternatives in this order? Nover and Hájek correctly object that linking order to ‘natural’ indexing in this way will not work. The notion of ‘naturalness’ is notoriously slippery, and different mathematical redescriptions of the probabilities and payoffs yield different natural ways of indexing the game. Thus, we may face a situation where a game, G , has no intrinsic ordering, where there are several different ways of describing G which yield different ‘natural’ indices, and where these different ways of indexing yield radically different expected values for playing G . If such games exist (or could exist) then Principle O is in trouble. At best O would be incomplete, failing to give guidance on how to value a game such as G .

Now, only games which yield a (countably) infinite number of possible outcomes can be sensitive to order in the above manner. So the key question is, are unordered, infinite games possible? For purposes of discussion I shall divide such games into *potentially* infinite games and *actually* infinite games. The PG is a potentially infinite game: there is no upper bound on the number of coin tosses needed to complete the game, but any given playing of the game will consist of some finite number of tosses. It seems clear that potentially infinite games must be ordered. Conceptually speaking, the very notion of a potentially infinite process presupposes an ordered sequence of substeps. Practically speaking, how can the mechanism inside the black box calculate whether or not to pay out for each positive integer n without taking the different n 's in *some* order?³ If this is right then an infinite game with no intrinsic order can only be an actually infinite game.

What if we allow for the possibility of actually infinite devices, for example a supermachine or superagent which can complete an infinite number of distinct operations in a finite time?⁴ This has the effect of dropping probabilities out of the discussion. Unlike the potentially infinite case, it seems possible to conceive of actually infinite yet unordered games. For example,

- (A) God gives you $1/n$ dollars for each positive integer, n , in an infinite number of simultaneous, instantaneous acts.

Were you to agree to play game (A), you would end up, seemingly, with an infinite amount of money. (The sum of transactions here is well-defined,

3. Note (and this is a point overlooked by Nover and Hájek when they describe the ‘black box’ scenario) that in the case of PG the separate outcomes are *not* independent. Exactly one of the infinitely many n -outcomes can occur.

4. Note that most discussions of ‘supertasks’ still involve an ordered sequence of infinitely many subtasks: for example a device which switches a light on or off, with the delay between successive switchings halving each time.

because the series $1/1, 1/2, 1/3, 1/4, 1/5, \dots$ is absolutely convergent.) But is game (A) itself well-defined? Compare the following game:

- (B) God gives you $1/m$ dollars for each odd positive integer, m , and takes away $1/n$ dollars for each even positive integer, n , in an infinite number of simultaneous, instantaneous acts.

What would you end up with after playing game (B)? Any arrangement of the transactions for game (B) forms a series which is conditionally convergent. Since the game itself is unordered, the sum is undefined. This suggests that game (B) is itself incoherent, and not something that even God could conceivably implement.

But how can God implement game (A) if He cannot implement game (B)? The individual acts involved in each game are of the same sort (i.e., giving or taking away small finite amounts of money) and the total number of acts in each is the same. I maintain—although this is not crucial to my overall position—that game (A) is also incoherent, even for an omnipotent God. This incoherence is masked by the absolute convergence of every series formed from the component acts. Nonetheless, these utilities must have *some* intrinsic order if they are to form a well-defined series. As game (A) is described they do not, hence game (A) is incoherent.⁵ Therefore all coherently specifiable infinite games are ordered in some way, and so Principle O applies to all such games.

4. The Intuitive Effects of Reordering. Even if—as I have argued—the apparatus of standard decision theory is sensitive to order, and even if all well-defined infinite games are ordered in some way, there remains a further issue about whether Principle O has intuitive plausibility. Does it make sense, in other words, that merely shifting the order in which possible outcomes might occur can make a difference to the intuitive value of a game?

The intuitive effects of switching ordering can perhaps be seen more easily in the case of actually infinite games. Consider the following—rather tedious—game, which I shall call the *Pass-a-Dollar Game* (PAD): in round 1, player A gives \$1 to player B; in round 2 player B gives \$1 to player A, and they continue alternating in this way for $2n$ rounds (where n is some positive integer). The sequence of transactions, viewed from player A's perspective, has utility $1, -1, 1, -1, \dots$. The overall expected value of the game is 0, since after an even number of rounds the partial

5. The same issue also arises in an epistemic version: God offers to pay you the sum of an infinite series consisting of infinitely many 1's and infinitely many -1 's in some order which you don't know. Should you take the bet? I want to argue that you have insufficient information on which to base a decision.

sum is always 0. Next consider the actually infinite version of the Pass-a-Dollar Game. A supermachine is constructed which can perform the transactions involved in PAD an infinite number of times in a finite period. At no point in infinite PAD can you have more than \$1 or less than \$0, hence the value of the game must presumably lie somewhere in this narrow range.

The expected value of infinite PAD is highly sensitive to ordering. Consider the following variation: we rearrange the series of transactions into a three-round cycle where you give the machine \$1, then it gives you \$1, then it gives you \$1 again, and then this is repeated *ad infinitum*. In formal terms all that has changed is the ordering of the terms in the series; now we have $-1, 1, 1, -1, 1, 1, \dots$. Yet it is clear that you make \$1 net profit after every three rounds of this game, hence the intuitive value of this game variant is way higher than infinite PAD.

So it seems that some actual infinite games are sensitive to the order in which the outcomes occur, even when this change in order has no effect on the value of any individual outcome. What about potentially infinite games? A detailed answer to this question is beyond the scope of this paper, but the above analysis already indicates ways in which order sensitivity may be a factor in such games also. From the perspective of decision theory, there is a formal equivalence between actually infinite games and potentially infinite games, since each kind of game generates an infinite series of expected utilities. So if an actually infinite game such as infinite PAD is order sensitive, why not a potentially infinite game such as PG? To argue against this possibility merely on the basis of its *prima facie* counterintuitiveness is not enough; in the realm of the infinite, intuitions are a notoriously poor guide to truth.

5. A New Paradox? I have argued that matching the order of the decision-theoretic formalism to the actual order of the situation it describes blocks the possibility of there being games to which decision theory is forced to assign multiple conflicting values. Thus PG, as described by Nover and Hájek, is not paradoxical. However, their other worry (which I have argued is less serious)—concerning the possibility of decision theory “remaining silent” on the value of playing certain games—remains. Given that we accept Principle O, is it possible to specify a game which is coherent, possible to play (i.e., only potentially infinite, not actually infinite), and yet decision theory fails to assign it any value?

My answer to this question is “yes.” As an example of a game with the above characteristics, consider the following, which I call the Enhanced Alternating St. Petersburg Game (EASP). A coin is tossed repeatedly until it comes up heads for the first time. The payoffs are as follows:

1st head on round 1: \$2
 1st head on round 2: -\$8
 1st head on round 3: \$24
 . . .

In general, the payoff for the first heads coming up in round n is $\$(-1)^n \cdot n \cdot 2^n$. The expected value is:

$$\begin{aligned} \text{EU}(\text{game}) &= 2 \cdot 1/2 - 8 \cdot (1/2)^2 + 24 \cdot (1/2)^3 - \dots \\ &= 1 - 2 + 3 - 4 + \dots \end{aligned}$$

The sequence of partial sums of this series diverges. This in itself is problematic, though no more so than the corresponding series for the standard St. Petersburg game which also diverges. However this latter series tends monotonically towards $+\infty$, allowing a unique (albeit counterintuitive) value of $+\infty$ to be calculated for the game. The real problem with the EASP sequence is that it *oscillates*. The series does not converge, nor is it bounded in any way.

My claim, then, is that EASP possesses the crucial problematic properties which Nover and Hájek attribute to PG, but that—unlike PG—EASP's problems are untouched by Principle O. Their worry stemmed from the alleged “silence” of decision theory concerning the value of PG. I argued—at the beginning of Section 2—that the real problem with PG is not underdetermination (incompleteness) but overdetermination (inconsistency). Decision theory seems to allow for multiple conflicting values to be assigned to PG, however this can be blocked by adopting the independently plausible Principle O.

EASP is a genuine example of the “silence” which originally bothered Nover and Hájek. It seems impossible to pin a unique value on EASP, nor even a value which is bounded within some finite range. Moreover, this indeterminacy of value seems to be independent of any constraints linked to the ordering of the game. Once our expectations are put in order, the Pasadena Paradox dissolves. However, the example of the Enhanced Alternating St. Petersburg Game suggests that their original worries were not unfounded: here is a game in which the order, utility, and probability of each outcome is fully specified and yet decision theory (even with the addition of Principle O) fails to assign it any value at all. What the correct response to this new paradox should be is far from clear.⁶

REFERENCE

Hájek, Alan, and Harris Nover (2004), “Vexing Expectations”, *Mind* 113: 237–249.

6. I am grateful to audience members at the Philosophy of Science Association Meeting in Vancouver, and to Alan Hájek, for helpful comments on earlier versions of this paper.