UNBOUNDED EXPECTATIONS AND THE SHOOTING ROOM

RANDALL G. MCCUTCHEON

Abstract. Several treatments of the Shooting Room Paradox have failed to recognize the crucial role played by its involving a number of players unbounded in expectation. We indicate Reflection violations and other vulnerabilities in extant proposals, then show that the paradox does not arise when the expected number of participants is finite; the Shooting Room thus takes its place in the growing list of puzzles that have been shown to require infinite expectation. Recognizing this fact, we conclude that prospects for a "straight solution" are dim.

1. Introduction

Several well-known puzzles and paradoxes philosophers have been discussing in the previous couple decades, including the the St. Peterburg Paradox (see e.g. Martin 2001), the Two Envelopes Problem (see e.g. Chalmers 1994), the so-called Pasadena Game (Nover and Hájek 2004), etc., crucially involve quantities unbounded in expectation. The Shooting Room (Leslie 1996) is a puzzle that clearly involves a quantity unbounded in expectation. Heretofore, however, published treatments have failed to even recognize this–much less indicate in a clear way that the involvement is crucial.

William Eckhardt (1997) reworks the puzzle to his comfort zone:

> Successive groups of individuals are brought into a room and given the same highly favorable wager, say, betting $100.00 that the "house," with fair dice, rolls anything but double sixes. (In the original formulation, losing players are shot, but this added gruesomeness, if nothing else, complicates the question of how one should bet.) Whenever the room occupants win their bets, ten times as many people are recruited for the next round. Once the house wins, the game series is over. So the house can truthfully announce before any games are played that, in spite of the highly favorable odds, at least 90% of all players will lose. The puzzle is that these bets appear to be both favorable and unfavorable: favorable because double sixes are rare, unfavorable because the overwhelming majority of players lose.

We will make a few cosmetic changes/addenda to the above scenario. First, one person bets in the first round, nine in the second round, ninety in the third, then nine hundred, nine thousand, etc. Next, the players are stipulated to have epistemically similar backgrounds and bet in isolation from each other; they can't tell how many others are betting in that round. Finally, the results of the bets are only announced at an open debriefing attended by all the players from all the rounds of betting. Players may count the number of attendees, but they sleep between playing and debriefing, and do not know how much time has passed in the interim.

Eckhardt argues for a credence of $\frac{1}{36}$ in *I lose* prior to the roll one is betting on, but is also comfortable, so far as we can tell, in a credence $\approx \frac{9}{10}$ in *I lose* once the bets are concluded:

> ...it is an error to consider yourself a random or typical player until you lose your bet. Before then, you have only about a 3% chance of belonging to the 90% majority. (...) This means *you should not consider yourself random until the game series is over.* (...) if a player about to bet were truly random among all players, then he would have better than a 90% chance of losing....

What troubles us is the emphasized (in the original) clause, which appears to sanction a belief, by the player, that she has lost with 90% probability once "the game series is over" (at the debriefing, in our version). No "solution" endorsing this can claim to have dissolved the paradox; a player having credence $\frac{1}{36}$ in *I lose* at the time the bet is placed, knowing that with probability 1 the game will end (to deny this leads to trouble too; see below) and she will have credence of at least $\frac{9}{10}$ at the debriefing, appears to violate Reflection. (And this does not appear to be one of the well-known "benign" sorts of violations of "naive" Reflection; cf. Schervish et. al. 2004).

In an inspired moment, Paul Bartha and Christopher Hitchcock (1999) offer a way out of the seeming paradox. Assigning "draft positions" to the potential participants that determine the order in which they will be called to the room, they show that if an individual has *any* countably additive probability distribution over her possible draft position, her expected credence in *I lose* at the debriefing will be precisely $\frac{1}{36}$.

To see this let p_n be the probability of a "draft round" equal to $n+1$ (e.g. if your "draft position" is in $\{11, 12, \dots, 100\}$, your "draft round" is 3) and note that:

(1) Prior probability in *I play* is $\sum_{n=0}^{\infty}(\frac{35}{36})^n p_n$.
(2) Posterior in *I lose* conditional on 10^x present at debriefing is $\frac{p_x}{p_0+p_1+\cdots+p_x}$.
(3) Probability of 10^x present conditional on *I play* is

$$\frac{Pr(10^x)Pr(play|10^x)}{Pr(play)} = \frac{(\frac{35}{36})^x(\frac{1}{36})(p_0+p_1+\cdots+p_x)}{\sum_{n=0}^{\infty}(\frac{35}{36})^n p_n}, \text{ hence}$$

(4) Expectation of posterior in *I lose* at debriefing conditional on *I play* is

$$\sum_{x=0}^{\infty} \frac{(\frac{35}{36})^x(\frac{1}{36})(p_0+p_1+\cdots+p_x)}{\sum_{n=0}^{\infty}(\frac{35}{36})^n p_n}\Big(\frac{p_x}{p_0+p_1+\cdots+p_x}\Big) = \frac{1}{36}.$$

This solves the Reflection problem but introduces a new one in its place. Bartha and Hitchcock put it this way: "...the weakness of this analysis is its inability to accomodate the intuition that (the participant) is equally likely to have any draft position."[1] They then offer another analysis in which participants have a merely finitely additive distribution over draft positions assigning equal infinitesimal weight to each natural number n.

Without going into the specifics of how, Bartha and Hitchcock conclude that, although conditional on the participant being selected to enter the room and the experiment having ended with a losing round, the probability that she lost is indeed $\frac{9}{10}$, it is however the case that conditional just on the participant being selected, the probability that she lost is $\frac{1}{36}$. Their explanation, roughly, as for why this is so is that, conditional on the participant being selected to enter the room, the probability that the experiment ends isn't equal to 1, but rather[2] $\frac{5}{162}$!

$$
\begin{aligned}
&Pr(Lose|Selected)\\
=&Pr(End|Selected)Pr(Lose|End \,\&\, Selected)\\
&\quad + Pr(No\ End|Selected)Pr(Lose|No\ End \,\&\, Selected)\\
=&\left(\frac{5}{162}\right)\left(\frac{9}{10}\right) + \left(\frac{157}{162}\right)\left(\frac{0}{1}\right) = \frac{1}{36}.
\end{aligned}
\tag{1.1}
$$

But this merely pushes the paradox back a round. Indeed, Bartha and Hitchcock look to be committed to a probability of the participant being in a final winning round, conditional on the participant being selected to enter the room, of $\frac{1}{360}$; she is one-tenth as likely to be in a final winning round as in a losing round. That is:

$$
\begin{aligned}
&Pr(Final\ Winning|Selected)\\
=&Pr(End|Selected)Pr(Final\ Winning|End \,\&\, Selected)\\
&\quad + Pr(No\ End|Selected)Pr(Final\ Winning|No\ End \,\&\, Selected)\\
=&\left(\frac{5}{162}\right)\left(\frac{9}{100}\right) + \left(\frac{157}{162}\right)\left(\frac{0}{1}\right) = \frac{1}{360}.
\end{aligned}
\tag{1.2}
$$

More generally, they should say that conditional on her entering the room, the probability that there are exactly n winning rounds after hers is $(\frac{1}{10})^{n+1}\frac{1}{36}$. Indeed, observe that $\sum_{n=-1}^{\infty}(\frac{1}{10})^{n+1}\frac{1}{36} = \frac{5}{162}$; apparently this is what Bartha and Hitchcock would assent to. Yet this is just as puzzling as the original paradox. To bring this into relief, stipulate that you only win if the rolls in your own round and the next land other than double six. (That is, it requires now two safe rolls to win.) Nothing else changes. Now it seems (by the Principal Principle, which Bartha and Hitchcock appear to defer to in their endnote 3) that you should assign probability $\frac{1}{36} + \frac{35}{36} \cdot \frac{1}{36} = \frac{71}{1296}$ to

[1] Since we are assuming that the players have identical epistemic backgrounds, presumably they will each employ the same distribution over draft position. The result is that they won't in general believe, at a debriefing attended by 10^n players, that their own draft position is, say, 1 with probability 10^{-n}. In other words, each player will believe themselves to be "special".

[2] We don't agree that the probability should be $\frac{5}{162}$, but it's more straightforward to grant this and demonstrate that it doesn't resolve the paradox anyway than it would be to recapitulate their sixteen page nonstandard analysis argument and say where we think it overreaches.

losing. On the other hand, it seems you should assign probability at least $\frac{99}{100}$ to losing (since at least that fraction of all participants lose). But no...according to Bartha and Hitchcock, the probability that you lose is actually only $\frac{1}{36} + \frac{1}{360} = \frac{11}{360}$! (Which, bizarrely, is now *less* than it would be if there weren't so many losing participants.)

We don't guess that *anyone* will be able to live with this. So until Bartha and Hitchcock offer an explanation for why one should preserve the accord with the Principal Principle for the first roll of the dice, but break it for subsequent rolls, one must conclude that the $\frac{5}{162}$ gimmick fails to accomplish anything at all.

2. Infinite repeated Shooting Room

At the beginning of the paper we promised an analysis of the Shooting Room paradox showcasing the role of infinite expectation. To this end it is instructive to contrast a *finite* expectation variant of the Shooting Room in which players lose their bets with probability $\frac{35}{36}$; e.g., they win on double sixes and lose otherwise. Rounds with identically mounting betting populations are again conducted until some group loses. Call the venue for these bets the *Brutal Room.* The Brutal Room game again requires an unbounded participant pool in order to ensure, with probability one, that one can complete the requisite series of rounds, and we'd again like to assign draft position priors in such a way that, at a postgame debriefing attended by N participants, each comes to have a draft position posterior that is uniform on $\{1, \ldots, N\}$.

As things are currently set up, however, this entails a uniform draft position prior over the infinite pool of potential participants. Merely finitely additive distributions (as we've just seen evidence of) are seldom, if ever, the way out of a paradox; more typically, they give rise to them. A more useful first step is to recognize that one *can* identify a countably additive draft position distribution that, when taken as our participant's distribution conditional on *I play*, yields a uniform debriefing posterior.

We identify the distribution. Denote by x the probability, conditional on *I play*, of draft round $D = 1$. Since 9 times as many enter with draft round $D = 2$ (namely positions 2-10), but round 2 takes place on only $\frac{1}{36}$ of iterations, the probability, again conditional on *I play*, of draft round $D = 2$ is $x(9)(1/36) = \frac{1}{4}x$. Similar reasoning shows the probability, conditional on *I play*, of draft round $D > 2$ is $\frac{1}{4}x(\frac{5}{18})^{D-2}$.

Setting the sum of these probabilities equal to 1 and solving for x we get:
$Pr(D = 1) = x = \frac{26}{35}$,
$Pr(D = 2) = \frac{13}{70}$,
$Pr(D = 3) = \frac{13}{252}$, etc.

If now a participant shows up at the debriefing and there are (for example) 10^3 participants present, her posterior draft position (not round, but position) distribution will be uniform on $\{1, 2, \ldots, 100\}$. This is because she knows one of (a)-(c) holds:
(a) $D = 1$ and two double sixes rolled after her arrival (probability $q_1 = \frac{26 \cdot 35}{35 \cdot 36 \cdot 36 \cdot 36}$)[3],
(b) $D = 2$ and one double six rolled after her arrival (probability $q_2 = \frac{13 \cdot 35}{70 \cdot 36 \cdot 36} = 9q_1$),

[3]To be clear, q_1 is the probability, for a participant at the debriefing who has not yet counted those present, of the conjunction "$D = 1$ and exactly 3 rounds were played".

(c) $D = 3$ and zero double sixes rolled after her arrival (probability $q_3 = \frac{13 \cdot 35}{252 \cdot 36} = 10q_2$).

This holds generally...regardless of how many players our participant encounters at the debriefing, she will come to have uniform draft position posterior over the corresponding initial segment of $\mathbf{N}$. Of course one might think that there is still a potential complication...to achieve the draft position distribution we found conditional on *I play* would still require a uniform prior, should we continue to assume that the game is played exactly once and that players are informed as to their status as pool members prior to learning that they have been selected to enter the room.

There is no such complication. Simply assume that the Brutal Room game has been played infinitely many times throughout an infinite past, and will continue to be played throughout an infinite future. (On the view that the universe is infinite, this strikes us as the most natural attitude.) Indeed, we may assume for convenience here that every being is chosen to participate in exactly one iteration of the game. (Perhaps they are selected in order of birth.) Once a given participant has commenced the game, the first participant in that particular iteration identifies an "origin" and the participant may revert to our talk of "draft position" and 'draft round". Since *I play* is now a certainty, however, there is no need for the problematic uniform prior.

A Brutal Room participant suffers no Reflection violation. During her betting session, her credence in *I will be alone at the debriefing*, is $Pr(D = 1)$ times the probability that the first roll is not double six, i.e. $(\frac{26}{35})(\frac{35}{36}) = \frac{13}{18}$. Credence in *I win* at the debriefing is 0 if they are alone and $\frac{1}{10}$ otherwise, with expectation $(\frac{5}{18})(\frac{1}{10}) = \frac{1}{36}$.

Buoyed by this result, one might think to analyze the Shooting Room along similar lines. We start by attempting to find a countably additive draft position distribution that, if taken as a Shooting Room participant's distribution conditional on *I play*, yields a uniform debriefing posterior. Denote then by x, as before, the probability, conditional on *I play*, of draft round $D = 1$. Since 9 times as many enter with draft round $D = 2$ (namely positions 2-10), and round 2 takes place on $\frac{35}{36}$ of iterations, the probability, conditional on *I play*, of draft round $D = 2$ is $x(9)(35/36) = \frac{35}{4}x$. Similarly the probability, conditional on *I play*, of draft round $D > 2$ is $\frac{35}{4}x(\frac{175}{18})^{D-2}$. But in order for these not to sum to ∞, one must have $x < \alpha$ for every positive real α, so our participant is again saddled with a merely finitely additive prior. A Reflection violation, too: the participant's credence in *I win* is $\frac{1}{36}$ during her betting session, while her credence in *I win* at debriefing is 0 if she is alone and $\frac{1}{10}$ otherwise, with expectation $(1 - x)\frac{1}{10} \approx \frac{1}{10}$.

More generally, if p is the (objective) probability that bettors win in a Shooting Room style game (with round sizes fixed at 1, 9, 90, etc.) then the (subjective) probability x of draft round $D = 1$, conditonal on *I play*, is equal to the multiplicative inverse of the expected number of bettors in a given iteration of the game. x is therefore infinitesimal (and the bettors' credences conditional on *I play* merely finitely additive) if and only if the expected number of players, namely $1 + 9p\sum_{n=0}^{\infty}(10p)^n$, is infinite (that is, precisely when $p \geq \frac{1}{10}$). Since Reflection violations don't arise when the participants have countably additive distributions, then, games of this sort require infinite expectation in order to generate a paradox.[4]

3. Two possible resolutions

To this point we've been content to steer discussion to the Shooting Room's crucial feature, namely the slow decay rate of $10^k Pr(X \geq 10^k)$, where X is the size of the participant pool. Given, in particular, that the paradox evaporates when $E(X)$ is assumed finite, the puzzle takes its place in a larger group of similar puzzles depending crucially on infinite expectation. At this point one sees two sorts of "resolutions".

Michael Huemer (2018) epitomizes the first. He claims that faithful implementations of the Shooting Room are impossible, writing "The paradox depends on metaphysically impossible assumptions about an infinite population of potential victims...."[5] He starts by noting a tension between the following two plausible principles.

> **Objective Chance Principle:** Given that event A happens if and only if B happens, and B has an objective chance c of occurring, the probability that A happens is c.
>
> **Proportion Principle:** Given that x is an A, and the proportion of A's that are B is p (with no reason for regarding x as being more or less likely to be B than any of the other A's), the probability that x is B is p.

The Objective Chance Principle pinpoints probability of winning at $\frac{35}{36}$; according to Huemer, the Proportion Principle points to a probability of $\frac{1}{10}$. Huemer now writes "The two principles yield the same result...for all possible population sizes. They disagree only in the impossible case of an infinite population." Several arguments are then given for this "impossibility".

Though this strategy is promising, Huemer's implementation fails to discern the role of expectation. It's not general enough to treat the population size X as a constant, nor even uniformly bounded: for every N there is *some* non-zero chance that the population size might exceed N. What the tension between the two principles turns

[4]The way we have set things up, a paradox can be generated (as in the Shooting Room scenario) whenever the size of the participant pool X is taken to be distributed in such a way that the sequence $10^k Pr(X \geq 10^k)$ grows at least exponentially; that is, when there is some $\beta > 1$ such that $10^k Pr(X \geq 10^k) \geq \beta^k$ for every (large enough) k. When $10^k Pr(X \geq 10^k) \leq \alpha^k$ eventually for some $\alpha < 1$, meanwhile, paradox is averted (as in the Brutal Room scenario)–we suspect that all nomologically possible scenarios are of this sort. There are of course intermediate cases.

[5]Elsewhere Huemer attempts a finitary reduction: "The impossibility...does not really lie in the fact that the scenario assumes an infinite population. (...) Since the game has a first round, it must have begun at some time. Whatever that time was, only a finite time has elapsed since then." Huemer then references a calculation in which it is shown, in effect, that if many separate game series were started N rounds ago then an expected $\frac{35}{36}$ of their participants up to now would have won (independently of N). A crucial feature of this calculation, however, is that these game series may not have ended, so it doesn't suggest any way at all to avoid the apparent Relection violation between time of play and *the end of the series*. (Cf. Eckhardt 1997, who is explicit that the 90% solution doesn't gain traction "until the game series is over", and Bartha and Hitchcock 1999, whose agent subscribing to the 90% solution explicitly "hears the news that the game has ended".)

on is the decay rate of $10^k Pr(X \geq 10^k)$. The conclusion Huemer should aim for is something like "The two principles yield the same result...for all possible distributions over the size of the participant pool. They disagree only for impossible distributions."[6]

David Chalmers (2002) is an instance of the second sort of resolution. Addressing the Two Envelopes Problem, Chalmers' strategy is to first formalize what he takes the argument to be that there is an expected benefit to switching envelopes, then offer a "disjunctive analysis", attacking one set of premises when the expected amount in the envelopes is finite and a different premise in the infinite expectation case. In particular, he calls attention to a new scenario that, to his mind, makes it more clear that the latter premise is defective in cases where infinite expectations arise.

Might one do something similar here? *Yes*, though the results don't, to our minds, resolve the paradox entirely.[7] First one would attack the Proportion Principle by creating a scenario in which it appears to fail. For a simple example, imagine an infinite row of houses, each containing an ideally rational agent. A fair coin is tossed in each and placed under a cup prior to observation. The agents are put to sleep and the houses rearranged so that the coin lies *heads* in every third house. Knowing all this, what should an agent's credence in *heads* be at wakeup? One half, or one-third?

One would argue (having chosen this path) that the answer is one-half. The permuting is after-the-fact and arbitrary: one could just as easily have arranged that every third house be *tails*, or any other pattern. Repeat the experiment many times and one's coin will have landed *heads* about half of the time with high probability, regardless of how the houses are reshuffled in one's sleep. So when told, at wakeup, that exactly one of three visible houses is a *heads* house, one should not invoke the Proportion Principle at all. Either the principle is false, or it doesn't apply[8]; if an agent meets with her two new next door neighbors on the lawn, knowing that these neighbors are in the same epistemic situation she is in and that exactly one of them has a *heads* coin, she should nevertheless consider herself to be "special" and continue to assign probability one-half to her coin lying *heads*.[9] And (one would reason), what

[6]We doubt infinite expectation distributions are possible, but Huemer's arguments, which assume constant, not random, pool size, can't purport to establish this in their current form.

[7]Chalmers' (2002) proposal for Two Envelopes (infinite expectation case) also breeds demons. The agent there, taking Envelopes A and B to confer infinite expected utility, declines to swap A for B though it would be in her interest to do so for every possible value in A. Fair enough, but consider a scenario where she is offered 1 util to swap, then will learn the amount in A, then again be given opportunity to swap (but without the 1 util enticement). She will decline the initial swap opportunity but accept the latter, leaving her 1 util worse off than if she had accepted the first. As always, admission of infinite expectation engenders vulnerability to a plethora of "bad books".

[8]It's perhaps impossible to craft a convincing counterexample to the principle in its current formulation. That would require a case where the probability that x is B isn't p, and yet where one has no reason to regard x as being more or less likely to be B than any of the other A's. If one's credences obey the probability axioms, such a case would exhibit the property that one has no reason to regard x as having the probability claimed for it. So the claim the argument is making is really that this (trivial) principle doesn't apply here. (Nor, by analogy, in the Shooting Room.)

[9]Again, the rationality under scrutiny isn't robust; these agents could be trivially subjected to a group Dutch Book, for example. Indeed, once one allows infinite expectations, we concede with Vann McGee (1999) that "rational decision making is a lost cause" (McGee's so-called "Airtight Dutch

goes for the coin problem goes for the Shooting Room. In the infinitely repeated version we considered, there are, for every $k \geq 0$, infinitely many groups of agents of size 10^k. Each group wins a bet with probability $\frac{1}{36}$ and they should hold on to that probability regardless of how they are sorted into "debriefing rooms" in their sleep. They might be sorted into rooms where half of all agents have won, or rooms where three or ninety-nine percent have won. Their attitudes to all such sortings should be the same–indifference. The actual game's protocol provides a way of sorting agents into rooms (where ten percent have won) that should be ignored like any other.

Our purpose here isn't to adjudicate between these rival approaches (ruling them out vs. salvaging some impoverished version of rationality in their wake) to infinite expectations. Each may have merit. We do insist, however, that the notion of infinite expectation ought to figure crucially in the analysis of any puzzle in which infinite expectation's role is in fact crucial. Analyses that fail this test (*all* previous analyses of the Shooting Room, in particular) have, we urge, missed the point–at least in part.

References

Bartha, Paul and Christopher Hitchcock. 1999. The Shooting-Room Paradox and Conditionalization on Measurable Challenged Sets. *Synthese* 118(3):403-437.

Chalmers, David. 1994. The Two-Envelope Paradox: A Complete Analysis? Manuscript. Available at http://consc.net/papers/envelope.html

Chalmers, David. 2002. The St. Petersburg Two-Envelope Paradox. *Analysis.* 62:155-57.

Eckhardt, William. 1997. A Shooting-Room View of Doomsday. *Journal of Philosophy* 94(5):244-259.

Huemer, Michael. 2018. *Paradox Lost.* Palgrave-McMillan.

Leslie, John. 1996. *The End of the World.* New York: Routledge.

Martin, R. 2001. The St. Petersburg Paradox. in E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy (Fall 2001 Edition)* [online encyclopedia], Fall 2001.

McGee, Vann. 1999. An airtight Dutch Book. *Analysis.* 59:257-265.

Nover, H. and A. Hájek. 2004. Vexing Expectations. *Mind* 113:237-249.

Ross, Jacob. 2010. Countable additivity, Dutch books, and the Sleeping Beauty problem. *Philosophical Review* 119:411-447.

Schervish, M.J., T. Seidenfeld and J.B. Kadane. 2004. Stopping to Reflect. *Journal of Philosophy* 101(6):315-322.

Department of Mathematical Sciences, University of Memphis
E-mail address: rmcctchn@memphis.edu

Book" involves a sequence of individually favorable payoffs X_i such that $\sum_{i=1}^{\infty} \min\{X_i, 0\} = -\infty$) and with Jacob Ross (2010) that "there are contexts in which full rationality is impossible" (Ross considers a Sleeping Beauty experiment with a number of awakenings having infinite expectation). Where we differ from these authors is in our recognition of infinite expectation's role.