Strong and Weak Expectations

Kenny Easwaran

Fine has shown that assigning any value to the Pasadena game is consistent with a certain standard set of axioms for decision theory. However, I suggest that it might be reasonable to believe that the value of an individual game is constrained by the long-run payout of repeated plays of the game. Although there is no value that repeated plays of the Pasadena game converges to in the standard strong sense, I show that there is a weaker sort of convergence it exhibits, and use this to define a notion of 'weak expectation' that can give values to the Pasadena game and many others, though not to all games that fail to have a strong expectation in the standard sense.

Terrence Fine shows that any valuation whatsoever of the Pasadena, Altadena, and St Petersburg games is compatible with the axioms of linear utility (which he also calls 'the axioms of rational preference', (Fine 2008, p. 618)). Thus, he seems to suggest that there are basically no normative constraints on one's preferences involving these games. However, one may wonder whether 'consistency expressed in the axioms' is the correct normative notion. It is conceivable that these axioms are either too weak (leaving out some normative constraints), too strong (imposing constraints where none are actually rationally required), or both.

As Fine points out (Fine 2008, pp. 627–628), the Archimedean requirement of linear utility contradicts one particular formulation of a principle of dominance. Since 'accepting dominance contradicts our fundamental assumption that utility is real-valued' (Fine 2008, p. 629), he does not think that 'acceptance of [dominance] is dictated by basic notions of rational preference'. (Fine 2008, p. 628) Others (like Nover and Hájek, or Colyvan (2006)) may naturally read this result in a manner more favourable to dominance than the Archimedean axiom (and thus the assumption that utility is always real-valued).

Be that as it may, there is still a question as to whether the axioms of linear utility are sufficient for rationality, or whether further rational requirements may be found that do put constraints on the values of the Pasadena, Altadena, and perhaps St Petersburg gambles. Fine suggests that there is not: 'We do not think that acceptance of [several conditions beyond the axioms of linear utility] is dictated by basic notions of

Mind, Vol. 117 . 467 . July 2008
doi:10.1093/mind/fzn053

rational preference' (Fine 2008, p. 628). However, I suggest that there might be such extra constraints.

But even if such constraints are not rationally required, Fine suggests that '[l]inear utility theory is about 'rationally' reducing the complexity of choices between gambles' so we can 'delegate to the mathematics' (Fine 2008, p. 620) rather than having to make decisions directly. If he is right about this, then there may be a pragmatic (if not a 'rational') reason to come up with extra constraints—they will let us delegate more decision problems to the mathematics.

1. Strong and weak expectations

As Fine points out, expectations play many roles. In the case of simple gambles (ones with only finitely many distinct outcomes), the axioms of linear utility require that an agent value the gamble at the expectation of the values of its component constant gambles. (Fine 2008, p. 624) Savage's decision theory (which allows weighted addition of gambles, rather than probabilistic mixture) may be more familiar to many philosophers than Fine's linear utility theory, but it too requires that agents value simple gambles at their expectations. However, even for non-simple gambles, the Strong Law of Large Numbers entails a further fact about expectations. If the X_i are independent, identically-distributed random variables whose expectation EX is given by an absolutely convergent sum, and we define $S_n = X_1 + \ldots + X_n$, then with probability 1, $\lim_{n\to\infty} S_n/n = \mathrm{E}X$ (Fine 2008, p. 616). That is, a long sequence of plays at any price above EX will almost certainly lead to a loss, and a long sequence of plays at any lower price will almost certainly lead to a gain.

Given that the goal is to find appropriate rational norms for decision making, one might wonder which (if any) of these three facts about expectations gives them normative force. Fine has shown that linear utility says nothing about preferences involving the Pasadena, Altadena, and St Petersburg games, and I suspect that Savage's decision theory may also be similar. (It cannot even discuss the St Petersburg game, rather than allowing it to be compared to other games.) But if the behaviour of the game in an infinite sequence of plays is the force behind expectation as a normative price for a single play, then we might look to see if anything like the Strong Law of Large Numbers applies in the cases of interest.

As it happens, for the Pasadena and Altadena games, something does. The Strong Law says that for certain distributions X, if we define a sequence of independent random variables X_i with this distribution, and define $S_n = X_1 + \ldots + X_n$, then for any ε,

$$P\left(\lim_{n\to\infty} |S_n/n - EX| < \varepsilon\right) = 1$$

If we let X be the Pasadena distribution, then this result does not hold for any value EX. In fact, as Fine states (and I prove in Appendix A), with probability 1 the sequence of S_n/n achieves both arbitrarily high and low values. Thus, if one player keeps getting to decide whether to play again or quit, then she can almost certainly guarantee as much profit as she wants, regardless of the (constant) price per play. However, as I will prove in Appendix B, for any ε,

$$\lim_{n\to\infty} P\left(|S_n/n - \log 2| < \varepsilon\right) = 1$$

That is, by fixing in advance a high enough number n of plays, the average payoff per play can be almost guaranteed to be arbitrarily close to log 2.

The difference between these two results is an interchange of the limit and the probability. Effectively, the first condition requires that in almost every sequence of plays, the average payoff eventually approaches EX. The second one only says that by fixing a large enough number of plays, we can make the probability of being arbitrarily close to EX be arbitrarily high. The former entails the latter, but the example of the Pasadena game shows that the converse is not true. As probability theorists put it, this is the distinction between 'almost sure convergence' in the first case and 'convergence in probability' in the second.

In general, if there is some value EX such that

$$\lim_{n\to\infty} P\left(|S_n/n - EX| < \varepsilon\right) = 1$$

then I will say that EX is the 'weak expectation' of the distribution used to generate the S_n. In the more standard case, where S_n/n converges to some value almost surely, I will call that value the 'strong expectation'. Because almost sure convergence entails convergence in probability, strong expectations are themselves weak expectations, but the example of the Pasadena game shows that the converse does not hold. Various versions of the Strong Law of Large Numbers give conditions under which the standard expected value is a strong expectation, while the Weak Law of Large Numbers gives more general conditions under which there is a weak expectation.

2. Which expectation (if either) matters?

This of course leads to an interesting question—is there any sense in which an agent is rationally required to value an individual gamble at its weak expectation? I will give an argument that might support such a claim, together with some responses and rejoinders, though I think the discussion is inconclusive. But if there is such a requirement, then Fine's result gives a criticism of linear utility theory, rather than an actual endorsement of the permissibility of any valuation of the Pasadena game. As Alan Hájek has pointed out in personal communication, this would suggest a programme to find a new set of axioms giving a sort of representation by weak expectations, preferably extending some standard set of axioms connected to strong expectations.

For the argument in favour of valuing games at their weak expectations, consider the case of an agent playing a game a fixed very large number of times. If she plays repeatedly at a price that is slightly higher than the weak expectation, then she has a very high probability of ending up behind. If she plays repeatedly at a price that is less than the weak expectation then she has a very high probability of ending up ahead. (Playing at a price exactly equal to the weak expectation for each play does not guarantee anything about the probability of eventually being likely to be ahead or behind. (Feller 1968, p. 249)) Because of this fact about repeated plays of the game, the agent ought to use the weak expectation as the guide to an *individual* play of the game as well—she should be willing to pay any amount less than the weak expectation to play, and should be willing to sell a play of the game for any amount more than the weak expectation.

Of course, one might wonder why it matters that this can generate a high probability of being ahead or behind—a single lottery ticket might be favourable despite giving one a very high probability of being behind. But one might think that in repeated play this has more normative force—if the lottery really is favourable, then repeated play will make up for the high probability of loss on individual plays. Also, by fixing the number of plays high enough at a price above the weak expectation, the probability of coming out ahead can be made arbitrarily small, while a single lottery play gives a fixed non-zero probability for coming out ahead.

Another worry about this argument is that the behaviour of games with weak but not strong expectations in repeated play becomes stranger if the number of plays is not fixed in advance. As mentioned earlier, with probability 1 the sequence of S_n/n for the Pasadena game goes arbitrarily far in both the positive and negative directions. Therefore, if the agent can choose when to stop playing, then she can be sure of being able to get ahead and then quit no matter what (constant) price she is paying per play. (This assumes that she has no limit to the debt she can sustain while pursuing this strategy, but it seems that this sort of assumption must already be at work for an agent to be able to play the Pasadena game even once.) On the other hand, if her opponent gets to choose when to quit, then she can be almost guaranteed to eventually come out behind, no matter how much she is being paid for each play.

But this response might not tell against the rationality of paying the weak expectation in normal situations, because a similar argument applies against any assignment of value to a simple coin flip. When betting at standard 1:1 odds on a fixed unfair coin coming up heads, an agent repeatedly using a 'double-or-nothing' strategy can ensure probability 1 of eventually coming out ahead, no matter how biased the coin is towards tails. Such a strategy to give a high probability of eventually coming out ahead in an unfair game is called a 'martingale'. We don't think the fact that the agent can almost guarantee herself to come out ahead with this sort of strategy makes the 1:1 odds fair here, if the bias is strongly towards tails.

Of course, in the case of Pasadena, the martingale is a repeated sequence of *identical* plays, while with the unfair coin she needs to keep doubling her bet, so there may be a disanalogy. Thus, this strange behaviour might be relevant, and the argument in favour of the weak expectation may fail.

A larger overall worry for the argument is that it claims to get a normative result about an individual play of a game from the behaviour of the game when it is repeated a very large number of times. (I thank an anonymous referee for pressing me on this point.) With strong expectations, the almost certain convergence applies even when we take the average of a sequence of distinct (but independent) random variables. (Feller 1968, Ch. X.5) Thus, even though an individual game may not be offered repeatedly, we can use the long-run behaviour of a sequence of distinct choices to argue that choosing the act with highest strong expectation at each point is the best strategy.

The referee has also pointed out a sort of alternative evaluation of games with no expectation (as suggested by Feller 1968, Ch. X.4). On the picture I have proposed, the value of a large sequence of plays of some game should just be a constant multiple of the number of plays, where the multiplier is the weak expectation of the game. On Feller's picture, the value of a large sequence of plays depends in some *non-linear* fashion on the number of plays. In particular, he suggests that a sequence of n plays of the St Petersburg game should have value $n \log n$. With this valuation, he demonstrates that something like the Weak Law of Large Numbers holds. A version of the Weak Law states that

$$\lim_{n\to\infty} \mathrm{P}\left(\left|\frac{S_n}{n\mathrm{E}X} - 1\right| < \varepsilon \right) = 1$$

For the St Petersburg game, we instead have

$$\lim_{n\to\infty} \mathrm{P}\left(\left|\frac{S_n}{n\log n} - 1\right| < \varepsilon \right) = 1$$

Thus, it is suggested that the longer a sequence of plays of the St Petersburg game, the more one should be willing to pay for each (one should be willing to pay $\log n$ for each of n total plays). This explains why no finite constant value is fair. In general, any two sequences that satisfy this property must be asymptotically equivalent, so any version of this result for Pasadena or Altadena must eventually be close to $n \log 2$ or $n(1 + \log 2)$, so this alternative approach only changes the valuation of games that otherwise get infinite value. This may suggest an important challenge to the additivity of value of independent gambles, but since we generally do not know how many choices of a given form we will eventually face, this seems to have no obvious relevance to the value of a single play. I am suggesting that long-run behaviour may be relevant to decisions in individual cases, but this alternate scheme gives values that depend on the specific length of the long run, so it cannot work in general cases.

An anonymous referee has pointed out that the weak expectation of a variable X can often be calculated as $\lim_{n\to\infty} \mathrm{E}X_n$ where X_n is a variable that is equal to X when $|X| \leq n$ but is 0 otherwise. Where the Pasadena game has a payoff of $(-1)^{i-1}2^i/i$ with probability 2^{-i}, consider another game where the payoffs are instead $(-1)^{i+1}2^i$. For this game, if n is above an odd power of 2 and below the next power of 2, $\mathrm{E}X_n = 1$, and $\mathrm{E}X_n = 0$ otherwise. Thus, this method of calculating the weak expectation does not work. While this alone does not entail that there is no weak expectation, I suspect that this is in fact the case—or at least if not for this game, it seems very likely that just as strong expectations fail to exist for many games, so do weak expectations. At any rate, there is a coherent, well-defined position available on which agents are rationally required to value gambles at a price equal to their weak expectations (in cases when they exist), even if this position is not fully supported. Thus, to support a position like Fine's, saying that every price is permissible, one should give arguments against weak expectations rather than just basing this permissiveness on some axioms. But although weak expectations may be the basis for a form of decision theory going beyond Fine's linear utility theory, they will not solve all the problems originally brought up in (Nover and Hájek 2004). There will still be some decision problems where these tools fail us.

USC School of Philosophy
Mudd Hall of Philosophy
3709 Trousdale Parkway
Los Angeles, CA 90089
USA
easwaran@usc.edu

KENNY EASWARAN

Appendices

These results can both be derived simply from theorems in Chapter VII.7 of [Feller, 1968] but I give proofs here to clarify what is going on.

Let X_n be the payoffs from a sequence of plays of the Pasadena game, so that the X_n are independent and identically distributed, with

$$P\left(X_n = (-1)^{i-1} \frac{2^i}{i} \right) = 2^{-i}$$

Let $S_n = X_1 + \ldots + X_n$. I will show that S_n/n converges in probability to log 2, but does not converge almost surely. In particular, $\forall \varepsilon \left(P\left(\left|\frac{S_n}{n} - \log 2\right| < \varepsilon \right) \to 1 \right)$ and $\forall \varepsilon\, P(|S_n/n| \geq \varepsilon$ for infinitely many $n) = 1$, which itself entails that there is no value to which S_n/n converges almost surely.

A. Lack of almost sure convergence

Claim: If $X_k \geq 2k\varepsilon$ then either $|\frac{S_{k-1}}{k-1}| > \varepsilon$ or $|\frac{S_k}{k}| > \varepsilon$.

Proof: If $X_k \geq 2k\varepsilon$ then $\frac{S_k}{k} \geq (S_{k-1} + 2k\varepsilon)/k = 2\varepsilon + S_{k-1}/k$.

If $\frac{S_{k-1}}{k-1} > -\varepsilon$ then we see that $S_k/k \geq 2\varepsilon - (k-1)\varepsilon/k > \varepsilon$.

Thus, if $X_k \geq 2k\varepsilon$ happens infinitely often, then $|S_n/n| \geq \varepsilon$ infinitely often.

Note that $\mathrm{P}(X_k \geq 2k\varepsilon) = \Sigma 2^{-i}$ where the sum ranges over all odd i such that $2^i/i \geq 2k\varepsilon$. By the Borel-Cantelli lemma, we see that if

$$\sum_{k=1}^{\infty} \mathrm{P}(X_k \geq 2k\varepsilon) = \sum_{k=1}^{\infty} \sum_{i \text{ odd}, \frac{2^i}{i} \geq 2k\varepsilon} 2^{-i}$$

is infinite, then with probability 1, $X_k \geq 2k\varepsilon$ occurs infinitely often.

Because all the terms in the sum are positive, we can rearrange the sum, to get

$$\sum_{i \text{ odd}} \sum_{k \leq \frac{2^i}{2i\varepsilon}} 2^{-i}$$

This sum is just

$$\sum_{i \text{ odd}} \frac{\lfloor 2^i/2i\varepsilon \rfloor}{2^i}$$

But $2^i/2i\varepsilon \geq \lfloor 2^i/2i\varepsilon \rfloor \geq 2^i/2i\varepsilon - 1$, so the sum we are interested in is bounded between $\sum_{i \text{ odd}} 1/2i\varepsilon$ and $\sum_{i \text{ odd}} \left(1/2i\varepsilon - 1/2^i\right)$. However, the difference between these sums is less than $\sum_{i=1}^{\infty} 2^{-i} = 1$ so the sum we are interested in is infinite iff either one of these sums is.

But $\sum_{i \text{ odd}} 1/2i\varepsilon$ is a harmonic series, and thus goes to infinity, so $X_k \geq 2k\varepsilon$ infinitely often with probability 1, and thus for any ε, with probability 1, $|S_n/n| \geq \varepsilon$ infinitely often, QED.

B. Convergence in probability

Now, I will show that S_n/n converges to log 2 in probability. To do so, I will use the Weak Law of Large Numbers:

> Let $X_1, X_2, \ldots$ be [independent, identically distributed random variables] with
>
> $xP(|X_i| > x) \to 0$ as $x \to \infty$
>
> Let $S_n = X_1 + \ldots + X_n$ and let $\mu_n = E\left(X_1 1_{(|X_1| \leq n)}\right)$. Then $S_n/n - \mu_n \to 0$ in probability. (Durrett, 2005, p. 41)

The notation for μ_n just means that μ_n is the expectation of the random variable that is 0 whenever $|X_1| > n$, and equal to X_1 otherwise. It is easy to see that in the case of the Pasadena game, for $2^k/k < n < 2^{k+1}/(k+1)$, we have $\mu_n = 1 - 1/2 + 1/3 - \ldots + (-1)^{n+1}/n$. Thus, $\mu_n \to \log 2$.

Thus, to show that the weak expectation of the Pasadena game is log 2, it will suffice to show that $xP(|X_i| > x) \to 0$. But again, when $2^k/k < x < 2^{k+1}/(k+1)$, we have $P(|X_i| > x) = 2^{-k}$. Thus, $xP(|X_i| > x) \leq 2/(k+1)$. But since $k \to \infty$ as $x \to \infty$, we see that this value goes to 0, so the Weak Law of Large Numbers applies, QED.

References

Colyvan, Mark 2006: 'No Expectations'. *Mind*, 115, pp. 695–702.

Durrett, Rick 2005: *Probability: Theory and Examples*. Belmont, CA: Thomson Brooks/Cole.

Feller, William 1968: *An Introduction to Probability Theory and its Applications*. New York: John Wiley & Sons.

Fine, Terrence L. 2008: 'Evaluating the Pasadena, Altadena, and St Petersburg Gambles. *Mind*, 117, pp. 613–632.

Nover, Harris and Alan Hájek 2004: 'Vexing Expectations'. *Mind*, 113, pp. 305–317.

