

The Bounded Strength of Weak Expectations

JAN SPRENGER AND REMCO HEESSEN

Abstract

The rational price of the Pasadena and Altadena Game, introduced by Nover and Hájek (2004), has been the subject of considerable discussion. Easwaran (2008) has suggested that *weak expectations*—the value to which the average payoffs converge in probability—can give the rational price of such games. We argue against the normative force of weak expectations in the standard framework. Furthermore, we propose to replace this framework by a bounded utility perspective: this shift renders the problem more realistic and accounts for the role of weak expectations. In particular, we demonstrate that in a bounded utility framework, all agents, even if they have different value functions and disagree on the price of an individual game, will finally agree on the rational price of a repeated, averaged game. Thus, we explain the intuitive appeal of weak expectations, while avoiding both trivialization of the original paradox and the drawbacks of previous approaches.

1 Introduction

Some probabilistic games have been remarkably refractory to a straightforward analysis in terms of Expected Utility Theory. The reason is that for these games, a straightforward expectation value does not exist. Most recently, the Pasadena Game has caught a lot of attention and has been discussed by a number of authors (Nover and Hájek 2004; Colyvan 2006; Easwaran 2008; Fine 2008). A fair coin is tossed repeatedly until it first comes up heads. Assume that this happens at toss n . If n is an odd number, the agent receives \$ $2^n/n$; if n is an even number, the agent has to pay \$ $2^n/n$. Is this game desirable or not, and what is its rational price?

Although the game itself is well defined, Expected Utility Theory does not assign it a definite value. Therefore it is hard to argue for a particular rational price. Even worse, it has been shown that in the case of the Pasadena Game and related gambles, a standard axiomatization of Expected Utility Theory allows for preferences that conflict with straightforward dominance reasoning (Fine 2008).

These problems apparently demand some extension of Expected Utility Theory. Easwaran (2008) proposes to value the Pasadena Game at its *weak expectation*, arguing that the average payoffs will, in the long run, be very close to this real number with arbitrarily high probability. This observation for a long series of Pasadena Games is then used to derive the rational price for a single game. However, neither Easwaran himself nor the subsequent discussion paper by Hájek and Nover (2008) dare a final verdict on the feasibility of the weak expectations approach. This paper discusses under which circumstances weak expectations have normative force.

After illustrating the problems of the Pasadena Game, we formulate Easwaran's proposal and point out that weak expectations cannot, by themselves, determine the rational price of the game (section 2). Our alternative proposal consists in wedding weak expectations to a more realistic bounded utility framework. Contrary to what one may be inclined to think, this move does not

trivialize the problem, but leads to an instructive convergence theorem and establishes the normative force of weak expectations for repeated games (section 3). Finally, we wrap up our results (section 4).

2 The Pasadena Game and weak expectations

The Pasadena Game has been introduced by Nover and Hájek (2004) as a challenge to Expected Utility Theory. The problematic nature of the game is due to some tricky mathematical properties. In order to illuminate them, and to keep a consistent terminology, we start off with two definitions:

Definition 1 A countable probabilistic game is a random variable $X : \Omega \rightarrow \mathbb{R}$, together with a measurable space $\Omega := \{\omega_j\}_{j \in \mathbb{Z}}$ (ω_j being the outcomes of the game indexed by the integers), and a probability measure P on Ω .

Obviously, the Pasadena Game qualifies as a countable probabilistic game (see table 1). The outcomes indicate at which coin toss ‘heads’ occurs first, and the function X maps those randomly selected outcomes to monetary payoffs that are identified with utility units (Nover and Hájek 2004, p. 239). Now, we define the *expectation* of a probabilistic game:

Definition 2 A countable probabilistic game X is (strongly) integrable if and only if

$$\sum_{j \in \mathbb{Z}} P(\omega_j) |X(\omega_j)| < \infty.$$

The value $\mathbb{E}[X] := \sum_{j \in \mathbb{Z}} P(\omega_j) X(\omega_j)$ is called its (strong) expectation¹.

The Pasadena Game is not strongly integrable because the series $\sum_{j \in \mathbb{Z}} P(\omega_j) |X(\omega_j)|$ takes infinite value. However, there are *arrangements* of the series $\sum_{j \in \mathbb{Z}} P(\omega_j) X(\omega_j)$ such that its value, interpreted as the expectation of the game, is finite. We get a problem of *arbitrariness*: depending on which arrangement of the terms in the series is chosen, it may converge to any value (examples are given in Nover and Hájek 2004).

Nover and Hájek convincingly argue that no particular arrangement, no particular order of summation is privileged. We do not have a reason to prefer a specific valuation of the game over another. On the other hand, they observe that the Pasadena Game is worse than, for instance, the *Altadena Game* (table 2) where the outcome probabilities are the same, but all payoffs are increased by \$1. By straightforward dominance reasoning, regardless of the rational price of the Pasadena Game, the Altadena Game is better. However, due to the failure

¹We may additionally allow for the strong expectation to be defined in cases where $\sum_{j \in \mathbb{Z}} P(\omega_j) |X(\omega_j)| = \infty$ but either $\sum_{j \in \mathbb{Z}} P(\omega_j) X(\omega_j) \mathbf{1}_{\{X(\omega_j) > 0\}}$ or $\sum_{j \in \mathbb{Z}} P(\omega_j) X(\omega_j) \mathbf{1}_{\{X(\omega_j) \leq 0\}}$ is finite. Since we do not consider any examples of this type, we set this case aside.

Outcome ω_j	...	ω_{-2}	ω_{-1}	ω_0	ω_1	ω_2	...
Probability $P(\omega_j)$...	1/16	1/4	1/2	1/8	1/32	...
Payoff (in \$) $X(\omega_j)$...	-4	-2	2	8/3	32/5	...

Table 1: The Pasadena Game.

Outcome ω_j	...	ω_{-2}	ω_{-1}	ω_0	ω_1	ω_2	...
Probability $P(\omega_j)$...	1/16	1/4	1/2	1/8	1/32	...
Payoff (in \$) $X(\omega_j)$...	-3	-1	3	11/3	37/5	...

Table 2: The Altadena Game.




Figure 1: A comparison of two random paths of the Pasadena Game (full line) and the Altadena Game (dashed line).

of strong integrability, the axioms of Expected Utility Theory do not determine preferences on those games so that we may value the Pasadena Game over the Altadena Game (Fine 2008). This is arguably weird.

This dilemma raises a debate about whether, and how, Expected Utility Theory should be extended as to account for dominance reasoning (Colyvan 2006, 2008; Easwaran 2009). Easwaran (2008) argues that even without such extensions, a study of the long-run behavior of the Altadena Game reveals that it is preferable to the Pasadena Game. These considerations are motivated by numerical simulations, too: the latter reveal that, as the number of repetitions increases, the average payoffs of the Pasadena Game are typically in a neighborhood of $\$ \log 2$, whereas the average payoffs of the Altadena Game are typically near $\$ 1 + \log 2$ (see figure 1).

One might take this to mean that the worries of Nover and Hájek (2004) are misplaced and we should simply take the expected value of the Pasadena game to be $\log 2$ after all, but this underestimates the peculiar behaviour of this game. The average payoff of the Pasadena game does not stabilize over time. Indeed, as Easwaran (2008, Appendix A) proves, with probability one the average payoff will eventually take both arbitrarily high and arbitrarily low values. On the other hand, if the game is often repeated, the probability that the average payoff at that time is far from $\$ \log 2$ will be very small. Easwaran makes this mathematically precise in the following way.

Let X_i , $i \in \mathbb{N}$, be independent realizations of the Pasadena Game, and let $S_n := \sum_{i=1}^n X_i$ be the sum of these games. Easwaran (2008, Appendix B) shows that the average outcome S_n/n of the Pasadena Game converges *in probability*

to a particular value which we call its ‘weak expectation’, namely $\log 2$. This suggests that convergence in probability—which we discuss in detail below—might be the crucial notion for determining the value of the Pasadena Game, as expressed in the following decision rule.

Weak Expectation Rule (WER): A countable probabilistic game X should be valued at its *weak expectation*.² This is the real number μ that satisfies, if $(X_n)_{n \in \mathbb{N}}$ denote i.i.d realizations of the game, and $S_n := \sum_{i=1}^n X_i$ is the sum of payoffs,

$$\forall \delta > 0 : \lim_{n \rightarrow \infty} P \left(\left| \frac{1}{n} S_n - \mu \right| \geq \delta \right) = 0. \quad (1)$$

Games that satisfy (1) for some $\mu \in \mathbb{R}$ are called *weakly integrable*. More specifically, the existence of a weak expectation μ means that for any tolerance margin δ and for a fixed number of plays that is large enough, we will, with probability $1 - \varepsilon$, end up with an average payoff that is close to μ :

$$\forall \varepsilon, \delta > 0 \exists N_0 \forall n \geq N_0 : P \left(\left| \frac{1}{n} S_n - \mu \right| \geq \delta \right) \leq \varepsilon. \quad (2)$$

This allows us to outplay an opponent whose price for the Pasadena Game deviates from $\log 2$, by choosing a sufficiently large number of plays N_0 for a given ε . Thus, it seems that $\log 2$ is the rational price of the repeated, averaged game: we should buy it at no higher price, and sell it at no lower price, to avoid to be taken advantage of.

Easwaran uses this convergence property to justify WER as a guide to an individual game. He explains the implications of (1) and (2):

If [an agent] plays [the game] repeatedly at a price that is slightly higher than the weak expectation, then she has a very high probability of ending up behind. [...] 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 as well. (Easwaran 2008, p. 636, original emphasis)

In other words, since we are dealing with independent, indistinguishable realizations of the very same game, the rational price of $\$ \log 2$ for the average payoff must be grounded in corresponding properties of an individual game. Thus, we should value a single Pasadena Game at $\$ \log 2$, a single Altadena Game at $\$ 1 + \log 2$, and our dominance heuristics are saved.

This motivation suggests that WER can guide us to the rational price of *all* weakly integrable games. This position is very attractive: First, WER rescues the dominance heuristics for the Pasadena and Altadena Games. Second, the weak expectation is always equal to the strong expectation if the latter exists.

However, it can be disputed whether the rationale underlying WER is sound. Assume that X is a weakly, but not strongly integrable game. WER states that with probability $1 - \varepsilon$, the payoffs of S_n/n will be within a close neighborhood of the weak expectation. Thus, equation (2) can only develop normative force

²The name ‘weak expectation’ stems from the fact that (1) and (2) occur in the Weak Law of Large Numbers.

if we are willing to *neglect* outcomes that occur with probability smaller than ε . Prima facie, this sounds reasonable. We can choose ε as small as we like. There seems to be some probability of success at which it is always rational to take a risk, even if some highly unpleasant outcome happens otherwise. We often ignore dangers that occur with very small chances. We cross the street even if there is a chance that some crazy car driver will kill us. We catch a flight on an airplane that might crash. And so on. Rational decision-makers are apparently justified to ignore outcomes that happen with arbitrarily small probability. In this case, it allows us to outplay an opponent whose game price deviates from the weak expectation μ , by choosing a sufficiently large number of plays n .

Now for our objection. Given ε , choose a suitable n for outplaying your opponent in the repeated, averaged game S_n/n . Let $M^- \in \mathbb{R}$ be such that

$$P(S_n/n < M^-) < \varepsilon. \quad (3)$$

Following the neglect rationale that underlies WER, we are now entitled to ignore all outcomes with payoff less than M^- . Evidently, the game that we obtain when neglecting those outcomes, for instance by setting them to zero payoff, is extremely desirable: the payoffs are bounded from below, but unbounded from above. Such a game must be valued at infinity since S_n/n fails to be strongly integrable.³ Contrarily, choose $M^+ \in \mathbb{R}$ such that

$$P(S_n/n > M^+) < \varepsilon. \quad (4)$$

Applying the neglect rationale once more would give us an extremely undesirable game. Thus, the neglect rationale leads to arbitrary valuations of a weakly, but not strongly integrable game, undermining the normative force of weak expectations. More generally, by eliminating payoffs in a set of arbitrarily small probability, the Pasadena and Altadena Games can be turned into unproblematic games that have an ordinary expected value.

Two objections to our argument deserve mention, but we think that we can reject both of them. First, we apparently neglect that WER builds on a *convergence* result that holds not only for a specific n , but for all $n \geq N_0$. However, all problems that we have pointed out for a fixed, but arbitrary n will transfer to any other $n \geq N_0$. So this objection does not take off the ground.

Second, the truncations in (3) and (4) appear to be biased in a particular direction, either towards very high or very low payoffs. The WER proponent could insinuate that we should adopt an *unbiased* procedure where very large and very low payoffs are discounted evenly. This attempt, however, amounts to a *petitio principii*: it makes clear that WER relies on a *privileged* way of neglecting certain outcomes. In fact, the neglect rationale yields arbitrary values as the rational price of a weakly but not strongly integrable game.

Thus, if we are licensed to ignore outcomes with arbitrarily small probabilities, the rational price of a game is in the eye of the beholder. Dependent on which improbable outcomes the agent decides to ignore, the game can be extremely good, extremely bad, or something in between. This awkward problem is analogous to the original dilemma of the Pasadena Game: in the same sense that there was no privileged order of summation (Nover and Hájek 2004), there

³Such a game is structurally similar to the St. Petersburg Game which will be introduced in the next section.

is no privileged negligible set of outcomes. WER fails to realize that the very nature of the Pasadena Game is wedded to what happens in outcomes with arbitrarily small probability. This is not surprising: S_n/n is weakly, but not strongly integrable,⁴ and different orders of summation yield different rational prices. Thus, the averaged games are still ill-behaved, and by no means less problematic than the original Pasadena Game. The convergence in probability just means that if we play often enough, the extreme payoffs will cancel each other out most of the time (i.e., with probability $1 - \varepsilon$). This increases the probability of landing near $\log 2$, but does not change the awkward general structure of the game.

All this has substantial implications for WER. The proponent of WER draws on the fact that for all tolerance margins δ , there is an N_0 such that for all $n \geq N_0$, S_n/n is close to the weak expectation with probability $1 - \varepsilon$. We have argued that leaving out some outcomes, however small their probability, is a completely arbitrary procedure which could also lead to very desirable and very undesirable games. Hence, the normative force of WER is severely undercut. We emphasize that our counterargument applies to all games that are weakly, but not strongly integrable, not only to the Pasadena and the Altadena Game.

Thus, do weak expectations have no prescriptive function for the valuation of weakly integrable games? They do if we make additional assumption. A specific and plausible assumption—bounded utility—allows us to show that weak expectations determine the rational price of a repeated, averaged game. We argue that such an assumption does not trivialize, but clarify the problem, and that it leads to fruitful and important insights.

3 A bounded utility perspective

A straightforward solution to paradoxes of infinite and indeterminate expectations that are due to unboundedly increasing payoffs is a bounded utility framework. There, each agent has a *value function* which maps monetary payoffs to utility or happiness units. Bounded utility implies that the agent's value function is bounded, that is, that there is a maximal amount of utility that money can confer, even if we have infinite amounts of money. This assumption rescues a lot of our intuitive judgments. For instance, most of us would not perceive a significant utility difference between \$ 10^{50} and \$ 10^{100} : for purchasing every good that we desire, there is no practical difference between both sums. This feature must not be confused with the claim that there is an amount of money where any additional money has no value. It only implies that if you are already super-rich, then the additional utility that you derive from extra money becomes negligibly small.

As shown by Fine (2008), Expected Utility Theory does not always determine preferences for games with unboundedly large payoffs. Bounded utility extends the scope of EUT to those games, and moreover, it is psychologically more realistic than unbounded utility. The latter stance would imply that a rational agent would exchange any very desirable outcome, such as long and happy life, for the following risky prospect: with an arbitrarily small chance (e.g., $1/10^{100}$), the agent has an even more positive experience, otherwise she loses everything

⁴This follows from S_n/n being the linear average of i.i.d. realizations of the Pasadena Game.

that she values. As observed by Aumann (1977), no rational agent would go for such a prospect.

Hence, it is not surprising that bounded utility explains some of our intuitions in paradoxical games. Take, for instance, the *St. Petersburg Game* where a fair coin is flipped, and where the player receives \$ 2 (4, 8, 16, ...) if the first ‘heads’ occurs on the first (second, third, fourth, ...) toss. According to Expected Utility Theory, the game has infinite expectation, but few of us would be willing to pay a substantial amount for playing such a game (Hacking 1980, p. 563). Most of us would not pay more than \$ 25 because in most cases, the payoff is close to zero. None of us would value the St. Petersburg Game over *any* finite payoff that we receive with probability 1. Bounded utility makes sense of this, because even extremely large monetary payoffs cannot confer more than a certain maximum amount of utility.⁵ This leads to a finite expectation, in agreement with our belief that the game is only finitely desirable. Thus, the St. Petersburg Game, and also the Pasadena and Altadena Games, become finitely desirable, and the paradox disappears.

These resolutions come at a price, however: apparent *trivialization* and *subjectivism*. If utility is bounded, all probabilistic games are strongly integrable. Games with unbounded payoffs lose their distinctive feature and behave like a standard game with bounded payoffs so that the (strong) expectation always exists. A bounded utility framework seems to miss the gist of the original paradox where the monetary payoffs were straightforwardly identified with utiles: ‘Identifying [...] dollar amount with utility, you naively compute the expected utility of the game as an infinite sum’ (Nover and Hájek 2004, p. 234).

This objection need not worry us too much: as argued above, it is questionable whether there is a consistent way to speak about rational decisions in games that have unbounded utility payoffs and a particular probabilistic structure. But we have to demonstrate that mapping unbounded payoffs to a finite range of utiles is not only a convenient way of dissolving the original paradox, but also a means of achieving substantial insights into the nature of weakly integrable games.

Moreover, on a bounded utility reading, different people typically assign different utilities to one and the same monetary payoff. Thus, the description of a probabilistic game (e.g., the Pasadena Game) does not determine an intersubjectively compelling rational price. Rather, the rational price of the game is in the eye of the beholder, and her individual value function. This subjectivist feature of bounded utility apparently prevents any interesting prescriptive claims.

Therefore, bounded utility, at first sight, does not seem to be a suitable candidate for resolving the paradox of the Pasadena Game, and assigning a rational price to that game. However, we will show that results in a bounded utility framework are far from being trivial or overly subjective: different agents with different value functions will finally agree on the rational price of a *repeated, averaged* Pasadena Game. This price is equal to the weak expectation. So we conclude that only in a bounded utility framework, the convergence property cited in WER develops the desired normative force. The setup below sets the

⁵This is not meant to imply that bounded utility is the only available solution to the puzzle posed by the St. Petersburg Game, but compared to other proposals, like assuming a logarithmic relationship between money and utility, it has the advantage of being generally applicable to games with unbounded payoffs.

stage for our positive results which are proved in the sequel.

Setup: Take a group of M agents $G = \{1, 2, 3, \dots, M\}$ with strictly monotonically increasing, bounded and continuous value functions $v_i : \mathbb{R} \rightarrow \mathbb{R}$, $i \in G$. Let $\|f\|_\infty := \sup_{x \in \mathbb{R}} |f(x)|$ denote the supremum norm. Then there is a common bound for the v_i :

$$C := \sup_{i \in G} \|v_i\|_\infty < \infty. \quad (5)$$

Taking \mathbb{E} as the expectation symbol, it is possible to prove a theorem that ensures a unique rational price in the long run (see appendix A):

Theorem 1 *Let S_n denote the payoff sum of n i.i.d. realizations of a weakly integrable game with weak expectation μ . Then, uniformly in $i \in G$:*

$$v_i^{-1} \left(\mathbb{E} \left[v_i \left(\frac{1}{n} S_n \right) \right] \right) \xrightarrow{n \rightarrow \infty} \mu. \quad (6)$$

Explanation of the Theorem: $\mathbb{E}[v_i(S_n/n)]$ denotes the utility agent i assigns to S_n/n (in keeping with Expected Utility Theory, this is the expected value of the value function). So $v_i^{-1}(\mathbb{E}[v_i(S_n/n)])$ can be interpreted as the price which agent i assigns to S_n/n . The theorem states that, as the number of games increases, each agent values the averaged game S_n/n at μ , and the differences between the individual valuations vanish. While there may be disagreement on the *absolute* amount of utility that S_n/n confers on the agents, they all agree that their personal amount of utility is *equal to the utility of obtaining μ* .

Corollary 1 *Theorem 1 applies to the Pasadena and Altadena Game. As n increases, the agents agree on a rational price of $\$ \log 2$ for the Pasadena Game and $\$ 1 + \log 2$ for the Altadena Game.*

Two features of the theorem should be mentioned: First, it assigns, despite the criticisms of the previous section, a substantial role to weak expectations. Easwaran's conjecture that $\log 2$ —the weak expectation—is the rational price of the Pasadena Game is vindicated for the repeated case in a bounded utility framework. Second, disagreement on the valuation of a single weakly integrable game is transformed into consensus on the valuation of the repeated, averaged game.⁶

This gap between repeated and single-case games sounds paradoxical at first sight, but we think that it is entirely reasonable. A valuation of an individual Pasadena Game depends to a very high degree on personal, subjective assessment of very large/very low payoffs. Therefore there are no reasons for calling disagreement on the rational price of a single game irrational. However, a repeated game converges, due to the existence of a weak expectation, in probability to some value. Due to bounded utility, the very large positive/negative payoffs fail to compensate for this effect. Therefore it is rational to agree on the average value of the Pasadena game in a long run of games, but it need not be rational to agree on the price of a single, individual game.

⁶This agreement result holds, a fortiori, also for strongly integrable games. Note further that the players do *not* agree on the rational price of the game S_n , because bounded utility is in general non-linear.

Finally we would like to mention two additional points. First, bounded utility, like any framework where utility increases with monetary payoffs, saves our everyday dominance heuristics so that the Altadena Game is always preferred to the Pasadena Game. Second, our results concern agreement on the assignment of monetary prices. It is arguably easier, and requires less additional assumptions, to experimentally confirm a result of that type, than it is to confirm a result that states agreement on the (possibly infinite) utility at which agents value such a game. Both features speak for the approach taken in this paper.

4 Conclusions

Our original claims can be divided into a negative and a positive part. The negative part (section 2) argues that the normative force of weak expectations in games akin to the Pasadena Game is undercut by an arbitrary neglect of certain outcomes. Since such an arbitrary procedure is required for a definite valuation of the game, weak expectations fail to establish a uniquely rational price.

Then, the positive part of our paper (section 3) argues that marrying weak expectations to bounded utility preserves the best of both worlds. Bounded utility is not only psychologically intuitive and helpful to understand actual decisions, it can also explain why the Altadena Game is preferred over the Pasadena Game. Crucially, it vindicates that subjects may disagree on the value of an individual Pasadena Game while they will eventually agree on the value of a repeated Pasadena Game. This price is equal to the weak expectation of the game (Theorem 1 and Corollary 1).

In other words, we may rationally disagree on the price of a weakly integrable game if we play it only once, whereas in the long run, we will have to agree on a rational price. Thus, we have discovered a sense in which weak expectations matter although a strong expectation does not exist, and clarified under which circumstances they can guide rational decisions.⁷

Tilburg Center for Logic and Philosophy of Science JAN SPRENGER
Tilburg University
P.O. Box 90153
5000 LE Tilburg
The Netherlands
j.sprenger@uvt.nl

Tilburg Center for Logic and Philosophy of Science REMCO HEESSEN
Tilburg University
P.O. Box 90153
5000 LE Tilburg
The Netherlands
remcoheesen@gmail.com

⁷We would like to thank Jason McKenzie Alexander, Ken Binmore, Peter Brössel, Mark Colyvan, Franz Dietrich, Kenny Easwaran, Alan Hájek, Stephan Hartmann, Pieter Ruys, and three anonymous referees of this journal for their valuable and inspiring comments on this paper. Research on this project was supported by the Veni Grant 016.104.079 ‘An Objective Guide for Public Policy’ by the Netherlands Organisation for Scientific Research.

A Proofs

Proof of Theorem 1: The proof proceeds in three steps: first, we show that it is sufficient to prove a claim different from theorem 1, second, we bound S_n/n from above, third, we prove that other claim with the help of the upper bound.

The convergence result of Theorem 1 states that for all $\Delta > 0$, there is a $N_0 \in \mathbb{N}$ such that for all $n \geq N_0$ and all $i \in G$:

$$\left| v_i^{-1} \left(\mathbb{E} \left[v_i \left(\frac{1}{n} S_n \right) \right] \right) - \mu \right| \leq \Delta. \quad (7)$$

Since the v_i are continuous and strictly monotonically increasing, they have continuous inverse functions v_i^{-1} . Then, there are $\tau_i > 0$ such that for all $x \in \mathbb{R}$ with $|x - v_i(\mu)| < \tau_i$, we have

$$|v_i^{-1}(x) - v_i^{-1}(v_i(\mu))| = |v_i^{-1}(x) - \mu| \leq \Delta. \quad (8)$$

Therefore, it suffices to show that $\mathbb{E}[v_i(S_n/n)]$ comes ‘sufficiently close’ to $v_i(\mu)$. In other words, we demonstrate that for $\tau := \min_{i \in G} \tau_i$, there is an N_0 such that for all $n \geq N_0$,

$$|\mathbb{E}[v_i(S_n/n)] - v_i(\mu)| \leq \tau. \quad (9)$$

Then, the theorem would be proven because (8) would imply that $|v_i^{-1}(\mathbb{E}[v_i(S_n/n)]) - \mu| \leq \Delta$ for all $i \in G$. It remains to show (9).

To do so, choose $\delta_0 > 0$ and $0 < \varepsilon_0 < 1$ such that

$$(1 - \varepsilon_0) \delta_0 + 2\varepsilon_0 C \leq \tau. \quad (10)$$

(Since C is finite due to (5), such numbers must exist.) Now, since S_n/n converges in probability, we know that

$$\forall \varepsilon, \delta > 0 \exists N_0 \forall n \geq N_0 : P \left(\left| \frac{1}{n} S_n - \mu \right| < \delta \right) \geq 1 - \varepsilon. \quad (11)$$

By continuity of the v_i , we can choose δ_i such that if $|S_n/n - \mu| \leq \delta_i$, then $|v_i(S_n/n) - v_i(\mu)| \leq \delta_0$. Setting $\varepsilon := \varepsilon_0$ and $\delta := \min_{i \in G} \delta_i$ in (11), we then obtain

$$\exists N_0 \forall n \geq N_0 \forall i \in G : P \left(\left| v_i \left(\frac{1}{n} S_n \right) - v_i(\mu) \right| < \delta_0 \right) \geq 1 - \varepsilon_0. \quad (12)$$

Now, we are in a position to complete the proof. Let n be at least as large as N_0 in (12). Then we apply straightforward dominance reasoning, bounding $|v_i(S_n/n) - v_i(\mu)|$ uniformly from above by a random variable that takes the value δ_0 with probability $1 - \varepsilon_0$, and $2C$ otherwise:

$$\begin{aligned} |\mathbb{E}[v_i(S_n/n)] - v_i(\mu)| &\leq \mathbb{E}[|v_i(S_n/n) - v_i(\mu)|] \leq (1 - \varepsilon_0) \delta_0 + \varepsilon_0 \cdot 2C \\ &\leq \tau, \end{aligned}$$

applying (10) in the last step. Thus, we have shown (9) and proven the theorem. \square

Proof of Corollary 1: Follows immediately from Theorem 1 because (i) the Pasadena and Altadena game satisfy the assumptions of Theorem 1, and (ii) their weak expectations are given by $\log 2$ and $1 + \log 2$, respectively. \square

References

- Aumann, Robert 1977: 'The St. Petersburg Paradox: A Discussion of Some Recent Comments'. *Journal of Economic Theory*, 14, pp. 443–5.
- Colyvan, Mark 2006: 'No Expectations'. *Mind*, 115, pp. 695–702.
- 2008: 'Relative Expectation Theory'. *The Journal of Philosophy*, 105, pp. 37–44.
- Easwaran, Kenny 2008: 'Strong and Weak Expectations'. *Mind*, 117, pp. 633–41.
- 2009: 'Dominance-based Decision Theory', unpublished manuscript, source: <http://www.ocf.berkeley.edu/~easwaran/papers/decision.pdf>.
- Fine, Terrence L. 2008: 'Evaluating the Pasadena, Altadena and St Petersburg Gambles'. *Mind*, 117, pp. 613–32.
- Hacking, Ian 1980: 'Strange Expectations'. *Philosophy of Science*, 47, pp. 562–7.
- Hájek, Alan and Harris Nover 2008: 'Complex Expectations'. *Mind*, 117, pp. 643–64.
- Nover, Harris and Alan Hájek 2004: 'Vexing Expectations'. *Mind*, 113, pp. 237–49.