KATHOLIEKE UNIVERSITEIT
LEUVEN

Faculty of Economics and
Applied Economics

Department of Economics

Value without absolute convergence
by

Luc LAUWERS
Peter VALLENTYNE

Econometrics

Center for Economic Studies
Discussions Paper Series (DPS) 11.04
http://www.econ.kuleuven.be/ces/discussionpapers/default.htm

March 2011

DISCUSSION
PAPER

Value without Absolute Convergence

Luc Lauwers, Center for Economic Studies, K.U. Leuven, Belgium

Peter Vallentyne, University of Missouri

Abstract: We address how the value of risky options should be assessed in the case where the sum of the probability-weighted payoffs is not absolutely convergent and thus dependent on the order in which the terms are summed (e.g., as in the Pasadena Paradox). We develop and partially defend a proposal according to which options should be evaluated on the basis of agreement among admissible (e.g., convex and quasi-symmetric) covering sequences of the constituents of value (i.e., probabilities and payoffs).

A finitely additive theory of (e.g., prudential or moral) value holds that, where there are only finitely many parts, the value of a whole is the sum of the value of its parts. We address the problem of how to extend this sum-principle when there are an infinite number of parts.

Sometimes the sum of an infinite number of values is a well-defined finite number. For example, the values 1, - ½, ¼, - 1/8, 1/16, …, $(-1/2)^n$, … add up to 2/3. Indeed, whatever the order with which the terms are added together, the resulting total is equal to 2/3. The series 1 - ½ + ¼ - 1/8 + 1/16 -…+ $(-1/2)^n$ +… is *absolutely convergent*, which means that it converges (i.e., has a finite limit) and the series of the absolute values of its terms also converges (1 + |-½| + ¼ + |-1/8| + 1/16 +…+ $|(-1/2)^n|$ +… = 2). For absolutely convergent series, rearranging the order of the terms does not change to resulting total, and thus the total of the terms is well-defined. Sometimes, however, the order does matter. For example, consider the values 1, -½, 1/3, -1/4, …, $(-1)^{n-1}/n$, …. When added up in the listed order, we obtain 1-½+1/3-1/4…+$(-1)^{n-1}/n$+… = log2 (approximately .69), but a different total shows up when added up in the following order: (1+1/3–½)+(1/5+1/7–¼)+(1/9+1/11–1/6)+ ⊙⊙⊙ = 1.5 log2 (approximately 1.04). These two series are *conditionally convergent*, which means that they converge but not absolutely. They do not converge absolutely, because the series of the absolute values of the terms, 1 + ½ + 1/3 + ¼ +…+ 1/n +… is infinitely large. There are many conditionally convergent sequences, but for illustration we shall focus on that of 1-1/2+1/3-1/4+…+$(-1)^{n-1}$/n+…, which is known as the alternating harmonic series. The relevant fact here is that, for conditional convergence (unlike the more well-behaved absolute convergence), different totals show up by manipulating the order in which the terms are added.

Given a set of values for which the total depends upon the order of addition, one might suppose that finitely additive value theories must remain silent about whether such an option is more valuable than some other option. After all, typically, there seems to be no reason to sum the terms in one order rather than in another. We shall argue, however, that, at least sometimes, not all orders of summing are admissible, and that when this is so, there can sometimes be determinate assessments of value.

Throughout, we focus solely on finitely additive theories of value (i.e., theories that, where there are finitely many parts, determine value of the whole by adding the values of the parts). Although we believe that much of what we defend can be extended to cover other kinds of theories of value, we shall not attempt to do so here.

1. Natural Structure

All standard approaches to value theory appeal to standard mathematical sums. As indicated in the introduction, these are well defined only if the sums are absolutely convergent (and thus not order-dependent). These approaches are silent about the assessment of options for which the sum is not absolutely convergent. We shall now argue that, where the locations of value have a “natural structure,” the value of an option is fully determined by summing component values (e.g., times) in an order that *respects the natural structure of the locations*. If all structure-respecting sums agree on the value of an option, then the option has that value, even if, for some possible order of summing a different total is obtained. Standard sums and absolute convergence, that is, are not necessary for the assessment of an option. The rest of this section will explain and defend this idea.

Suppose that the prudential value of a world (with no uncertainty involved), for an individual, is determined by the sum of the local values at times. Suppose that (1) time is discrete (i.e., for which between any two times are there are only a finite number of other times; as opposed to dense) and unbounded (extends indefinitely) in both directions, and (2) people come into existence and then live forever. Consider a world in which a given person has the following distribution of value at times: <1,1,-1,1,1,-1…>. Is this better than 0 units of value at each of those times?

Unlike the example in the introduction, the sum of the values in <1,1,-1,1,1,-1…>, in the given order, does not converge (have a finite limit), since that sum is infinite. It does, however, have a limit in the extended sense (which we shall use) that includes positive and negative infinity as possible values.[1] Moreover, like the case of conditional convergence, the sum of 1,1,-1,1,1,-1,… is *order-dependent*. In the given order, it sums to positive infinity, but in the rearranged order 1,-1,-1,1,-1,-1,1,… the limit is negative infinity.[2] We shall use the example of 1,1,-1,1,1,-1,… to introduce an approach to normative evaluation that will be extended in the next section to cover cases of conditional convergence. In this section, we shall argue that, even though 1,1,-1,1,1,-1,… has no order-independent sum, it is better than 0 at each time.

Times, unlike people, have a natural structure. First, they have a *natural order* in the sense that the notions “before” and “after” are well defined. Jan. 2, 2009, for example, is after Jan. 1, 2009 and before Jan. 3, 2009. Throughout, we shall assume for illustration that time is Newtonian and thus that there is a determinate frame-invariant natural order. Our basic points, however, are unaffected if the natural order is less determinate (e.g., in relativistic times). Second, there is a *natural distance metric* in the sense that there are facts about how far a given time is from another. (The natural distance metric entails, of course, a natural order, but some of the issues below depend only on the natural order.) If, for a given choice situation (e.g., decision under certainty), times are the only relevant basic locations of value (for a person), it is appropriate, we claim, for value theory to appeal to the natural temporal structure when summing local values. (This is more fully defended, in a different context, in Vallentyne and Kagan 1997.) Thus, although the standard sum of the numbers in the set {1,1,-1,1,1,-1,…} is not well defined, we claim that, *for the normative purposes of prudential evaluation*, only sums that respect the structure of the basic locations of value (e.g., times) are relevant. We shall develop this idea below and claim that the above world is (infinitely) prudentially better than 0 on each of the listed days. (Of course, one might deny that temporal locations are basic locations of prudential value, but in this section we shall introduce our general approach on the assumption that they are.)

When the set of basic locations of value has a natural structure (as times do), the value of an option, we claim, is determined by the limits of its total values in certain kinds of expanding sequences of "bounded" sets of locations that "cover" all the locations. This is understood as follows. Let a *bounded* set of locations be (1) a finite set, if the locations have no natural distance metric (e.g., for people), and (2) a set for which there exists a finite upper bound for the distance between any two members, if there is a natural distance metric (e.g., for times). Thus, for example, neither an infinite set of people, nor the infinite set of times extending from a given time infinitely into the future, is bounded. The infinite set of dense times (instants) between one o'clock and two o'clock on a given day, however, is bounded. For a set of basic locations (e.g., times), L, a *covering sequence of sets* $<S_1, S_2, \ldots S_n, \ldots>$ is defined as follows:

(a) For each i, S_i is a bounded subset of L.

(b) For each i, S_i is a subset S_{i+1}.

(c) For each member, l, of L there is an i such that l is a member of S_i.

This just says that $<S_1, S_2, \ldots, S_n, \ldots>$ is a sequence of expanding bounded subsets of locations, with each location included beyond some point in the sequence. For example, if the locations are times, and time starts at time 0 and extends indefinitely into the future, then the following is a covering sequence, where "[n,m]" denotes the (interval) set of points inclusively between n and m: <[0,1], [0,2], [0,3], ….>. Note that covering sequences can start with any set of locations. There is no privileged starting point.

We claim that the assessment of an option is determined by the limits of the values of the option relative to various covering sequences—as long as the option is such that any bounded set of locations contains only a finite total value. Call an option *boundedly finite* just in case it satisfies this condition. An option containing one unit of value at each integral unit of time and no value elsewhere is boundedly finite. If time is dense, then an option containing one unit of value at each moment of time is not boundedly finite (since there will be an infinite total between any two times). Obviously, assessing options that are not boundedly finite is far more complex than assessing options that are. In what follows, we limit our focus to the assessment of the latter.

Let us say that a covering sequence of sets is *admissible* just in case it satisfies certain additional requirements. We shall address possible additional criteria for admissibility below, but let us first note our basic claim (a dominance condition, the spirit and name of which come from the social choice literature on infinite utility streams; e.g., Von Weizsäcker (1965) and Atsumi (1965)):

Weak Generalized Catching-Up: For any two boundedly finite options, A and B:

(1) If, (a) *for each* admissible covering sequence, the limit of the sequence of the corresponding values

of A is *at least as great* as the limit of the corresponding values of B, and (b) *for at least one* admissible covering sequence, the limit of the sequence of the corresponding values of A is *greater* than the limit of the corresponding values of B, then A is *more valuable* than B.

(2) If, for *each* admissible covering sequence, the limit of the sequence of the corresponding values of A and the limit of the corresponding values of B are *equal and finite*, then A is *equally valuable* with B.

This says the following for a boundedly finite option, A: If, for each admissible covering sequence, the limit of the sequence of the corresponding values of A is equal to the real number v, then A has (determinate) value v. If the admissible covering sequences generate different limits, then A's value is *indeterminate*. If the different limits include both negative and positive infinity, then A's value is *radically* indeterminate. Otherwise, it is *partially* indeterminate.

For simplicity, we shall write as if a given option has a limit (finite or infinite) relative to each infinite covering sequence, but this is not so. Our more general theory appeals instead to the weaker notion of cluster points (in the set of extended real numbers, which include positive and negative infinity), which all infinite sequences have. In a later section, we define cluster points and reformulate the above principle accordingly. The key points, however, can be made more simply on the assumption that the limits exist.

What, then, are the substantive criteria of admissibility of covering sequences? Where times are basic locations, there are, we claim, two criteria for admissibility. One is that, because time has a natural order, the sets in a covering sequence must respect that natural order. This requires that the sets must be *convex* in the sense that if t_1 and t_2 are members of a given set, then so must all intermediate times. For example, one is not allowed to add the value at time 2 and the value at time 4 without adding the value at time 3. If time is the only dimension of basic locations of value, this is equivalent to requiring that the sets be *intervals*. This requirement rules out gimmicky sets that involve holes and gaps. This requirement makes sense because you can’t realize the value at two distinct times without realizing the values at all intermediate times.

Before addressing the second criterion of admissibility for covering sequences, we can note already how this approach assesses <1,1,-1,1,1,-1,…> above. Given that covering sequences must be convex, and hence intervals (in this one-dimensional case), we know that this option is infinitely valuable and hence more valuable than <0,0,0,0,...>. No matter what initial bounded interval one starts with, any interval expansion that involves at least five (adjacent) integral times will have a positive total (since there be more 1s than -1s). Moreover, as the interval is expanded, the sum of the values of <1,1,-1,1,1,-1,….>, relative to the expansion, goes to positive infinity (that is, increases without limit). Thus, the limit of the sequence of these values is infinite. (This result, of course, will remain valid, even if there are further restrictions on admissible covering sequences.) The above approach thus says that this option is infinitely valuable, which is better than the zero option. This, we claim, is the right assessment. Even though the set {1,1,-1,1,1,-1,…} has no well defined total the temporally ordered sequence <1,1,-1,1,1,-1,…> is, we claim, infinitely valuable.

We now allow time to be unbounded in both directions and we argue that a second criterion of admissibility is needed. Suppose, for illustration, that a given person is alive at each time and thus has always been alive and always will be. (This is admittedly a strange example, but we use it to illustrate a point that will be relevant when we turn to risky options in the next section.) Consider an option that has the following values for this person each day: <…-1,-1,-1,-1,2,2,2,2,2….>. We claim that this option has positive value (indeed infinitely positive value). This, of course, is highly controversial, and we won’t attempt to defend that claim here. Instead, we shall simply use it to motivate a second criterion of admissibility. Without a further condition, the following covering sequence of convex sets, with their associated values under the above option, would be admissible: <-1,-1,2>, <-1-1,-1,-1,-1,2,2>, <-1,-1,-1,-1,-1,-1,-1-1,2,2,2>, … (This expands by adding three -1s to the left for each 2 to the right.) At the limit, this sequence has infinitely negative value. This covering sequence of convex sets, however, expands more to the left than it does to the right. Given that time has a natural distance metric, this seems inadmissible.

A natural thought is that, if the set of locations is symmetric in a sense to be defined below, then the sets in covering sequences must also suitably symmetric. More exactly, let us say that a set of locations is *symmetric with respect to some point* (not necessarily in the set) just in case there is a point, p, such that, for any location in the set, the mirror image of this location with respect to p is also in the set. For example, the set of discrete moments {1, 2, 5} is not symmetric with respect to any point, whereas the set {1, 3, 5} is symmetric with respect to moment 3 and {1, 2, 6, 7} is symmetric with respect to moment 4. Furthermore, the set {0, 1, 2, …} of temporal locations with a beginning but no end has no point of symmetry; and the set {… -2, -1, 0, 1, 2, …} of temporal locations without beginning or end is symmetric with respect to *every point* (location). (Throughout, we only consider convex sets of locations.)

We claim that a second plausible condition on admissibility for covering sequences is that they be *quasi-symmetric*. Roughly speaking, this requires that, if the space of locations is symmetric in a certain way, then each set in a covering sequence must respect this symmetry. In the next section, we will give a more careful general characterization of this condition, but here we shall simply state the one-dimensional version: if the set of locations is symmetric with respect to some point, then, for a given covering sequence, there must be a point relative to which each set in the sequence is symmetric relative to this point. No restriction is imposed in the case where the set of locations is not symmetric with respect to any point. Thus, if time has a beginning but no end, there is no point of symmetry, and hence all sets satisfy quasi-symmetry. If, however, time is unbounded in both directions, then every point is a point of symmetry, and, *for a given covering sequence*, there must be a point relative to which all sets in the sequence are symmetric. Different sequences can be symmetric relative to different points, but all sets in a given sequence must be symmetric relative to a fixed point. Quasi-symmetry rules out gimmicky expansions in one direction rather than another where the space of locations is symmetric with respect to some point. It imposes no restrictions where the space is not symmetric with respect to some point.

Let us now reconsider the option above that has the following values each day for a given person: <…,-1,-1,-1,-1,2,2,2,2,2,…>. We noted that the sequence of values associated with covering sequence <-1,-1,2>, <-1-1,-1,-1,-1,2,2>, <-1,-1,-1,-1,-1,-1,-1-1,2,2,2>, … has a limit of negative infinity, whereas the option seems to have positive value. In this example, time is unbounded in both directions and hence is symmetric with respect to every location. Quasi-symmetry thus requires that, for a given covering sequence, there be a fixed location relative to which all sets in the sequence are symmetric. The above covering sequence violates this condition. The point of symmetry for each set shifts two positions to the left with each expansion and there is no fixed point of symmetry for all the sets. (Keep in mind that symmetry is a property of the set of locations, and not of the set of attached values.) If, however, we consider only covering sequences that are convex and quasi-symmetric, then, no matter what convex set one starts with, later sets must expand an equal distance in both directions in order to have the same point of symmetry. The result is that the total value in the covering sequence eventually becomes unboundedly positive (since eventually one location with value 2 will be added for each location added location with value -1). This, we claim, is the correct answer.

For brevity of expression, we shall say that one set is a *quasi-symmetric expansion* of another just in case (*i*) the former is a proper superset of the latter, and (*ii*), if the space of locations is symmetric, then both sets are symmetric with respect to the same point (which also is a point of symmetry of the space of locations). Each non-initial set in a quasi-symmetric covering sequence is thus a quasi-symmetric expansion of its predecessor.

We claim that the following two criteria for *admissibility* of covering sequences are plausible:

1. If the basic locations have a *natural order*, then all the sets in covering sequences must be *convex* (i.e., include all locations between any two included locations).
2. If the basic locations have, for given dimension, a *natural distance metric*, then the covering sequence must be *quasi-symmetric*.

Given these two criteria of admissibility, the temporal sequence, <…,-1,-1,-1,-1,2,2,2,2,2,...> is assessed as infinitely valuable. It's worth noting, however, that this approach does not always give determinate assessments. Consider the temporal sequence <…,-1-1-1,1,1,1,1,...>. The above approach says that its value is radically indeterminate. This is because the limit of the sequence of values for this option, relative to a convex quasi-symmetric covering sequence, depends on the starting point of the sequence. If one starts with the right-most -1, then quasi-symmetric expansions (e.g., adding the same number -1s as 1s) will have a limit of -1, but if ones starts with the second right-most -1, then quasi-symmetric expansions have a limit of -3 (since -1 will be added to each side before any 1s are added with the offsetting -1s). Indeed, by starting sufficiently to the left, or to the right, the limit can be made as large (positive), or small (negative), as one wants. Thus, the above approach does not always lead to determinate results and sometimes leads radical indeterminacy. This, we claim, is appropriate.

So far, we have introduced this approach without any appeal to conditional convergence. Let us now examine its implications for such cases. If the basic locations have no natural structure (e.g., they are people, as opposed to times), then all covering sequences are admissible (since there is neither natural order nor a distance metric that needs to be respected). In this case, an option (world) with merely conditionally convergent values will have, in full agreement with the standard mathematical assessment, a completely indeterminate value. This result is known as the Riemann rearrangement theorem: if a series is conditionally convergent, then one can obtain any given value (finite or infinite) as its limit by rearranging the order of addition.

Consider now distributions over time. Here there is a natural order, a distance metric. Let us suppose that the locations extend unboundedly in both directions. Thus, both convexity and quasi-symmetry are required covering sequences to be admissible. Consider the following distribution of value day by day: < …, -1/6, -1/4, -1/2, 1, 1/3, 1/5, …>. In assessing this option, the following is an important relevant fact (proved in the appendix):

Limit Agreement 1: All quasi-symmetric convex covering sequences of a given option have the same limit when all the following conditions holds: (1) the basic locations have a natural order and a natural distance metric, (2) there is only one dimension for basic locations (e.g., just time), (3) the option is boundedly finite, and (4) for any location, the values in other locations converge to zero as their distance from the given location tends to infinity.

Under the above conditions, the limit of a given option is the same no matter what the starting set is and no matter what the exact manner of convex and quasi-symmetric expansion is. Thus, we can determine the value of an option by calculating its limit value for any one of the admissible covering sequences. For < …, -1/6, -1/4, -1/2, 1, 1/3, 1/5, …>, we can thus start with the location containing 1 and take the limit as we expand quasi-symmetrically. This gives us the sum 1-1/2+1/3 …, which is equal to log2. Although an option with a different ordering of the same values (e.g., <…-1/10, -1/6, -1/2, 1, 1/3, -1/4, 1/5, 1/7, -1/8, 1/9,…>) can have a different limit, for the fixed ordering of a given option, the starting point and manner of quasi-symmetric convex expansion does not matter.

We believe that the above two criteria of admissibility for covering sequences are plausible. If basic locations have a natural order, then convexity rules out gimmicky orders of summing that involve holes or gaps. If the basic locations have a distance metric and are unbounded in both directions, then quasi-symmetry rules out gimmicky orders of summing that proceed more in one direction than another. It's important to keep in mind that our claim here is not a claim about mathematics. For many of the cases that we address, there is no mathematically well-defined sum. Our claim is a normative claim about the (prudential or moral) assessment of options. The claim is that the assessment does not depend merely on the mathematical sum of the unordered set of values; it also depends on the natural structure of the normatively basic locations of those values.[3]

We now note that the following two principles must be rejected by any approach, such as ours, that goes beyond standard mathematical sums, and that satisfies a standard dominance principle (i.e., if, for each location, one option is at least as good in all locations, and better in some locations, than another option, then that first option is better than the second):

Indifference for infinitely many indifferent changes: The result of applying an infinite number of indifferent changes to an option produces an option that is equally good with the original option.

Improvements for infinitely many improving changes: The result of applying an infinite number of improvements to an option produces an option that is better than the original option.

To see that an infinite number of *indifferent changes* can produce a worse option, consider the following sequences of options:

<1,1,1,1,...>

<0,2,1,1,1,...>

<0,0,3,1,1,1,...>

<0,0,0,4,1,1,1,1,...>

…

<0,0,0,…>

Each step in the sequence just shifts a finite amount of value one position to the right. On our approach, and almost any plausible finitely additive theory, such shifts produce an equally valuable outcome. Note, however, that, at the limit, an infinite number of such shifts produces <0,0,0,…>, which is dominated by each of the options in the sequence. (Here and below, we assume, for illustration, that convergence is relative to the product topology.) Hence, an infinite number of indifferent changes need not produce an indifferent outcome. This is indeed strange, but once one moves beyond standard sums, it's clear that one needs to reject certain principles that are unproblematic when standards sums apply.

To see that an infinite number of *improvements* need not produce an improvement, consider the following:

<1,1,1,1,1,...>

<0,3,1,1,1,…>

<0,0,5,1,1,1,…>

<0,0,0,7,1,1,1,...>

…

<0,0,0,…>

Each step in the sequence just shifts a finite amount of value one position to the right and then adds 1. On our approach, and almost any plausible finitely additive theory, such shifts produce an improvement (because of the added 1 along with the shift). Note, however, that, at the limit, an infinite number of such shifts produces <0,0,0,…>, which is dominated by each of the options in the sequence. Hence, an infinite number of improvements need not produce an improvement. Again, this is strange but unavoidable once one moves beyond standard mathematical sums. (For more discussion of this issue, see Lauwers and Vallentyne 2004).[4]

These incompatibilities are well known in the social choice literature on infinite utility streams (see, for example, Fleurbaey and Michel, 2003). We believe that dominance is sufficiently compelling to warrant rejecting these principles in the infinite case.

We shall now use the admissible covering sequence approach to assess risky options.

2. Assessing Risky Options

We shall understand risky options, for a single person, to be probability distributions over payoffs in units of value. Throughout we restrict our attention to discrete payoff variables for which probability is countably additive (i.e., for which the probability of a union of a countable number of disjoint events is the sum of their individual probabilities). For an arbitrary enumeration of the payoff values with a non-zero probability of being realized, we can describe an option by a countable set $\{\langle p_1, v_1\rangle, \langle p_2, v_2\rangle, \ldots, \langle p_n, v_n\rangle, \ldots\}$, where p_i is the (positive) probability with which the payoff v_i is realized and $p_1+p_2+\ldots+p_n+\ldots = 1$. In this notation, the expected value of a risky option is defined as the probability-weighted sum $p_1v_1 + p_2v_2 + \ldots + p_nv_n + \ldots$, if this sum is absolutely convergent (and thus the total does not depend on the order of summing).

Consider now the *Pasadena game* introduced by Nover and Hájek (2004).[5] A fair coin is flipped until a heads comes up, and one wins something if the number of flips is odd and loses something if the number of flips is even. More precisely, the payoffs, along with the associated probability are described by the following probability distribution: $\{<1/2, 2/1>, <1/4, -4/2>, <1/8, 8/3>, \ldots, <1/2^n, -(-2)^n/n>, \ldots.\}$. This gives us the following set of probability-weighted values: $\{1, -1/2, 1/3, \ldots, (-1)^{n+1}/n, \ldots\}$, which is conditionally convergent. The Pasadena game thus has no well-defined expected value; its probability-weighted sum depends upon the order of summation.

Does the appeal to essential temporal structure help assess this option? It depends on whether the payoffs have a temporal structure. If the coin is flipped once each day, with the payoff following a second later, then the payoffs have a temporal structure, and all structure-respecting (i.e., convex and quasi-symmetric) covering sequences would assess the option as having value log2, as indicated in the previous section.[6] The payoffs, however, need not have a temporal structure. Suppose, for example, that the first coin flip occurs now, the second flip in half an hour, the third flip 15 minutes later, and the interval between successive flips is reduced by half each time. Thus, all coin flips will take place within the next hour. Suppose further that the payoff is given two hours from now. Thus, all possible payoffs take place at one fixed time (e.g., in two hours) and there is no temporal order to which to appeal.

In what follows, we shall focus on versions of the Pasadena game, and other games with risky options with conditionally convergent value, in which there is no variable temporal dimension to the payoffs. We shall claim that there is nonetheless a relevant structure that must be respected in summing values, and that, for many conditionally convergent games, this produces a determinate answer about whether such a game is better than zero value.

The natural and therefore relevant structure comes, we claim, from the constituents in the summed values: they consist of a probability multiplied by a value. For example, the Pasadena game is not fully represented by the set {1, -1/2, 1/3, -1/4, …}. It has more structure than that. A fuller representation is the (probability distribution) set {<1/2, 2>, <1/4, -2/1>, <1/8, 8/3>, <1/16, -16/4>, …}, where it's understood that the first element of each ordered pair is a probability, the second is a payoff, and the two elements are to be multiplied together.

Where the summed values have constituents, we claim, admissible covering sequences must respect the natural structure of the constituents. More exactly, we propose that the structure-respecting approach of the previous section can be adapted as follows. Let L (the set of locations of possible constituents of value) be the set of all possible ordered pairs, <p,v>, where p is a number inclusively between 0 and 1 (the "probability"), and v is any real number ("the payoff"). A risky option O, is a countable subset of L, $\{<p_1,v_1>, <p_2,v_2>, \ldots <p_k,v_k>\ldots\}$ for which the p_i sum to 1. We shall continue to write, metaphorically, of locations of value, but strictly speaking these are *possible constituents* of value—<p,v> pairs—that may be part of a given option. A given option may include, or not, a given <p,v> pair.

Given that there is a distance metric for both the probability dimension and the payoff dimension of L, a *bounded* set of locations is a set for which the maximum distance—in probability and in payoff—between any two members is finite. Given that L itself is bounded (between 0 and 1) in the probability dimension, the boundedness of a set of locations is simply a matter of its boundedness in the dimensions of payoff.

In the spirit of the previous section, we stipulate that a covering sequence is admissible only if each set in the sequence is convex and the sequence is quasi symmetric. A set is *convex* just in case, if two points are included in the set, then so are all intermediate points (i.e., all convex combinations). Thus, if <.1, 8> and <.7, 2> are included so must be <.3, 6>, <.4, 5> and many others. Again, convexity rules out gimmicky holes and gaps. For the two dimensional case, to define quasi-symmetry of a covering sequence, we first need to define the number of dimensions (0,1, or 2) with respect to which the space of locations is symmetric. Say that a set of locations is *m-dimensional symmetric* (for m = 0, 1, or 2) just in case the maximal number, up to 2, of intersecting lines relative to which the set is symmetric is m.[7] (A set is symmetric with respect to a line just in case, for any point in the set, its mirror image with respect to the line is also in the set. See diagram for illustration.) In the two dimensional framework, if the set of locations is m-dimensional symmetric, then a covering sequence is *quasi-symmetric* just in case, if m ≥ 1 (i.e., there is symmetry in at least one dimension), there are m (i.e., 1 or 2) intersecting lines (where a single line is deemed to intersect itself) of symmetry for the space of locations relative to each of which each set in the sequence is symmetric.[8] For two-dimensional space, if both dimensions are unbounded in both directions, then the space is 2-dimensional symmetric, and any two intersecting lines are lines of symmetry for the space. In this case, quasi-symmetry requires that there is a pair of such lines such that all sets in the sequence are symmetric with respect to both these lines. (As a consequence, all the sets have the same center.) If, however, one dimension is unbounded in both directions, and the other dimension is bounded in both directions (which is the case that we shall focus on below), then the space of locations is again 2-dimensional symmetric, but there is only one line of symmetry in the doubly bounded dimension: the line parallel to the two boundaries midway between the two. Nonetheless, there are still an infinite number of pairs of intersecting lines of symmetry, since each straight line orthogonal to the two borders is a line of symmetry. In this case, quasi-symmetry requires that there be some pair of intersecting lines of symmetry relative to which all sets in the sequence are symmetric.

We believe that quasi-symmetry is a plausible condition of admissibility for covering sequences, since it rules out gimmicky expansions that go more in one direction than another in dimensions for which the space of locations is itself symmetric.

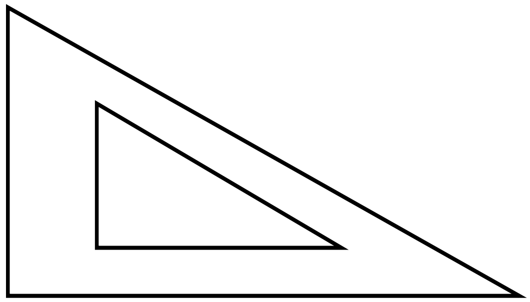

Two convex sets. There is no line of symmetry for either.

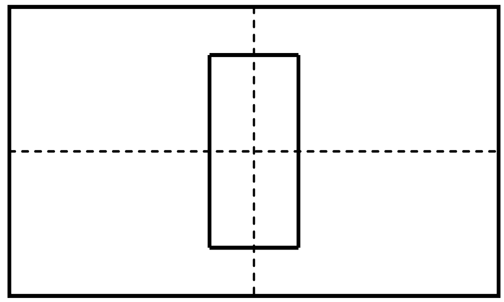

Two convex and 2-dimensional symmetric sets. The dashed lines are the lines of symmetry.

Let us now apply this approach to the 2-dimensional case of risky options. Here, probability is bounded in both directions (by 0 and 1) and payoffs are unbounded in both directions. Hence, the .5 probability line is the only line of symmetry for the probability dimension. Given the unbounded payoff dimension, each line orthogonal to the .5 probability line is a line of symmetry in the payoff dimension. Hence, the space is 2-dimensional symmetric. Thus, an admissible covering sequence consists of convex sets for which there is a pair of lines of symmetry (the .5 line and some orthogonal line) relative to which each set is symmetric.

Using these convexity and quasi-symmetry conditions for admissibility for covering sequences, we can now invoke the Weak Generalized Catching-Up principle (from above). As before, we restrict our attention to options that are *boundedly finite* (having a finite value relative to any bounded subset of L).

Weak Generalized Catching-Up (repeated)**:** For any two boundedly finite options, A and B:

(1) If, (a) for *each* admissible covering sequence, the limit of the sequence of the corresponding values of A is *at least as great* as the limit of the corresponding values of B, and (b) for *at least one* admissible

covering sequence, the limit of the sequence of the corresponding values of A is *greater* than the limit of the corresponding values of B, then A is *more valuable* than B.

(2) If, for *each* admissible covering sequence, the limit of the sequence of the corresponding values of A and the limit of the corresponding values of B are *equal and finite*, then A is *equally valuable* with B.

Admissibility: A covering sequence is admissible only if (1) if the locations have a natural order, then each set in the sequence is convex, and (2) if the locations have a natural distance metric, then the sequence is quasi symmetric.

We can now note the following:

The value of the Pasadena game: Each convex and quasi-symmetric covering sequence assigns the value log2 to the Pasadena game.

To see this, start by considering one specific convex, symmetric, covering sequence, which expands based on *the absolute value of the payoffs*. Consider <S_1, S_2, S_3, …, S_i,…>, where S_i is the set of all <p,v> pairs in L for which absolute value of v is less than or equal to $2^i/i$. Each set S_i is symmetric with respect to the .5 probability and the zero payoff-line. The total value of the Pasadena option, P, relative to this sequence, that is, <P(S_1), P(S_2), P(S_3),… P(S_i),…>, is the limit of the sequence 1 –1/2 + 1/3 + …+ $(-1)^{k+1}/k$ as k goes to infinity. As we know, this limit is equal to log2. This, of course, is only one of the infinitely many convex, quasi-symmetric covering sequences.

What are the limits for other kinds of convex, quasi-symmetric covering sequences for the Pasadena option? It turns out that, although this is not generally true, for the Pasadena option, all convex, quasi-symmetric covering sequences have the same limit. The argument is as follows. Consider an arbitrary convex and quasi-symmetric covering sequence. The lines of symmetry are the .5-probability line and a line orthogonal positioned at some value, v. Since we consider the limiting behavior, we may assume that the first set in the sequence already includes the points (1/2, 2), $(1/2^m, 2^m/m)$, and $(1/2^n, -2^n/n)$ with m and n sufficiently large (so that, due to symmetry with respect to the value-v-line and convexity, the points $(1/2^m,2)$ and $(1/2^n,2)$ also are included). Now, consider a set S further in the covering sequence and let $(1/2^k, 2^k/k)$ be the point in S with the highest (positive) associated value (1/k). (Thus, k is odd and greater than or equal to m.) Due to quasi-symmetry and convexity, the points $(1/2^k, 2)$, $(1-1/2^k, 2)$, and $(1-1/2^k, 2^k/k)$ also are included in S. Convexity implies that S includes the rectangle spanned by $(1-1/2^k, 2)$, $(1/2^k, 2)$, $(1/2^k, 2^k/k)$, and $(1-1/2^k, 2)$. Hence, all positive probability-weighted payoffs up to 1/k are taken into account. A similar argument establishes that, for some even natural number s, all negative probability-weighted payoffs that are at least as great as -1/s are taken into account. The sum of all these terms converges to log2 as t and s simultaneously go to infinity. Thus, for the Pasadena game, the limit of the sum of the probability-weighted payoffs is the same relative to all convex quasi-symmetric covering sequences.

We can generalize this argument by focusing on two features of the Pasadena game, {<1/2, 2/1>, <1/4, -4/2>, <1/8, 8/3>, …}: (1) the ordering of the pairs on the basis of decreasing probability is the same, with finitely many exceptions, as the ordering of the pairs on the basis of increasing absolute value of payoff (the only exception concerns the first two pairs); (2) the total probability mass in the positive and negative tails of the distribution is sufficiently small. The next condition clarifies the meaning of this second feature.

Thin Tails: For each positive real number r, the limit of the sum of probability-weighted absolute payoffs, $p_k|v_k|$, for $x \leq |v_k| \leq x+r$, goes to zero as x goes to infinity.

In the Pasadena game, this is satisfied. Recall that each (large) absolute value $|v_k|$ in the Pasadena game is almost the double of its predecessor $|v_{k-1}|$. Thus, for each positive real, r, for sufficiently large x, there is at most one $\langle p_k,v_k \rangle$ for which $|v_k|$ is between x and x+r. Moreover, given that the absolute value of the probability-weighted payoffs, $p_k|v_k|$, converge to 0 as x goes to infinity (1, ½, 1/3, 1/4 , …), it follows immediately that the sum of $p_k|v_k|$, for $x \leq |v_k| \leq x+r$, goes to zero as x goes to infinity. Thus, the Pasadena game satisfies Thin Tails.

We can now note (with the proof in the Appendix):

Limit Agreement 2: Let $O = \{\langle p_1,v_1 \rangle,\langle p_2,v_2 \rangle,\ldots,\langle p_n,v_n \rangle,\ldots\}$ be a risky option, enumerated on the basis of non-decreasing *absolute values* ($|v_n| \leq |v_{n+1}|$, for each n). Suppose that the ordering of these pairs according to non-increasing *probabilities* is the same, perhaps with finitely many exceptions, as the ordering listed above, and that $p_1v_1+ p_2v_2+\ldots+ p_nv_n+\ldots$ converges to V (finite or infinite). If Thin Tails is satisfied, then the limits of the sum of the probability-weighted values of this option are the same for all convex and quasi-symmetric covering sequences. In addition, if V is a finite number, then all these sums are the same *if and only if* condition Thin Tails holds.

Under the conditions of Limit Agreement 2, Weak Generalized Catching-Up assigns a determinate value to options. It is important note, however, that it does not always assign a determinate value to conditionally convergent options. To see how a disagreement in limits can arise, consider a variation on the Pasadena game in which the probabilities of the positive payoffs are subdivided and spread over nearby payoffs in a certain way but the probabilities of the negative payoffs are not changed.

Let us explain the changes with reference to the following table.

	Modified Pasadena Game	
Negative payoffs	Original Positive Payoffs	Revised Positive Payoffs
¼ x -4/2	1 location: ½ x 2	2 locations: ¼ x 2*
1/16 x -16/4		
1/64 x -64/6	1 location: 1/8 x 8/3	8 locations: 1/64 x 8/3*
1/256 x -256/8		
1/1024 x -1024/10		
1/4096 x-4096/12	1 location: 1/32 x 32/5	128 locations: 1/4096 x 32/5*
….		

The idea above is to subdivide the probabilities for the positive payoffs of the Pasadena game so that they are equal to the increasingly "far apart" probabilities for negative payoffs. The probability for the smallest positive payoff (1/2) is divided into two so as to be equal to the probability (1/4) of the largest negative payoff (-16/4). The probability for the *second* smallest positive payoff (1/8) is divided into eight parts so as to be equal to the probability (1/64) of the *third* largest negative payoff (-64/6). The probability for the *third* smallest positive payoff (1/32) is divided into 128 parts so as to be equal to the probability (1/4096) of the *fifth* largest negative payoff (-4096/12). And so on. For each subdivision of the probabilities, the associated payoffs are kept close to the original payoff (e.g., less than .1 difference) and are evenly balanced above and below the original value so that the net effect on probability-weighted value is zero (e.g., ½ x 2 becomes ¼ x 1.99 plus ¼ x 2.01). The above table references these evenly balanced numbers close to the original payoff using the asterisk sign as shorthand. For example, 2 x ¼ x 2* is shorthand for a payoff slightly below 2 and a balancing payoff slightly above 2, each with probability 1/4.

Let us now see that this option has a partially indeterminate value. Consider, once again, the admissible sequence <S_1, S_2, S_3,…, S_i,…>, where S_i is the set of all <p_j, v_j> pairs in L for which absolute value of v_j is less than or equal to j. Because the payoffs for this modified Pasadena game are only slightly different from the original Pasadena game, it is easy to show that, relative to this particular covering sequence, the above modified Pasadena game has the value log2. Consider now, however, a different admissible covering sequence. Consider <S^*_1, S^*_2, S^*_3, …, S^*_i,…>, where S^*_i is the set of all <p_j,v_j> pairs in L for which $1/j \leq p_j \leq 1 - 1/j$ and the absolute value of v_j is less than j. Like the previous sequence, each set is symmetric with respect to both the .5 probability line and the zero payoff line and hence quasi-symmetric. Unlike the previous sequence, however, there is a restriction on the probabilities in addition to the restriction on the payoffs. The total value of the modified Pasadena option, P*, relative to the sets in this sequence, is as follows: P*(S^*_1) through P*(S^*_3) are 0, since P* includes no points with probability greater than or equal to 1/3. P*(S_4) through P*(S_{63}) are each ½ [= 2x(1/4 x 2*) + (1/4 x -4/2)], since the three probability ¼ pairs are the only pairs with probability greater or equal to than 1/63. P*(S_{64}) through P*(S_{4096}) are each 5/12, since the thirteen pairs with probability greater or equal 1/4095 are 2x(1/4 x 2*) + (1/4 x -4/2) + (1/16 x -16/4) + 8x(1/64 x 8/3*) + (1/64 x -64/6) (= 1 -1/2 -1/4 +1/3 – 1/6). Splitting up the probability-payoff couples with a positive value into an increasing number of couples with smaller probabilities and slightly different values has the effect of progressively shifting the positive values to later in the sequence. As a result, the value becomes 1-1/2-1/4+1/3-1/6-1/8+… = .5 x log2.

Thus, for this modified option, the value is partially indeterminate (since both log2 and .5xlog2 are admissible values). It is not, however, completely indeterminate. Let us use this modified option to illustrate a method of determining the range of values that limits of admissible coverings sequences can take. In order to determine the *smallest* limit of an admissible sequence, we want to consider an admissible covering sequence that gives as much priority as possible to the *early inclusion* of probability-value couples with *negative* probability-weighted values. For illustration, let <1,0> (i.e., probability 1, payoff 0) be the starting set, S_0. Let S_1 be the smallest convex and quasi-symmetric set that includes <1/4,-2> (the probability-value couple with the lowest probability-weighted payoff) and <1,0>. This is the symmetric (with respect to the .5-probability line and the 0-payoff line) rectangle spanned by the four points <1/4,-2>, <3/4,-2>, <1/4,2>, and <3/4,2>. Let S_2 be the smallest convex and quasi-symmetric expansion that includes <1/16,-4> (the probability-value couple with the second-lowest probability-weighted payoff) and <1,0>. This is the rectangle spanned by <1/16,-4>, <15/16,-4>, <1/16,4>, and <15/16,4>. Observe that S_2 includes S_1. S_3 is the rectangle spanned by <1/64,-64/6>, <63/64,-64/6>, <1/64,64/6>, and <63/64,64/6>, and it includes S_2. And so on. The value of the modified game according to the sequence <S_1,S_2,…,S_k,…> is equal to .5 x log2 (as indicated above). This is the minimum admissible value of the modified game.

In order to find the *largest* limit on an admissible sequence, one uses a similar method but giving priority to probability-value couples with larger (rather than smaller) probability-weighted values. The maximal value of the modified game is equal to log2.

Thus, in the Pasadena game, our proposed approach yields a determinate value of log2, but in other games, such as the above modified Pasadena game, our approach yields partially indeterminate results (value is at least .5xlog2 and no greater than log2). In other games, it yields completely indeterminate results. It is, we claim, appropriately sensitive to the structure of such games.

It is, of course, strange that the modifications made above to the Pasadena game could somehow affect the value of the game. After all, each change made (splitting the probabilities) left the game equally valuable. How could an infinite number of such changes alter the value from log2 to a more indeterminate result? As noted in the previous section, however, in the context of an infinite number of locations of value, once one accepts a dominance (Pareto) principle and the principle that total-preserving *finite* shifts produce a result that is indifferent with the original, one must also accept than an infinite number of indifferent changes can change the evaluation of an option.

Here is one more example, based on moral evaluation. Suppose that there are an infinite number of people listed in some arbitrary order, and that their payoffs are <…,1,-1,1-1,1,-1,…>. Pick an arbitrary person with a payoff of 1 and shift her 1 one place to the right. She and her neighbor now each have 0. Let us suppose, for illustration, that the result is morally indifferent to the original, since payoffs were only finitely shifted and not increased in any way. Do this an infinite number of times in a manner that ensures that each (positive) 1 is shifted to the right exactly one position. The net result is <…0,0,0,…>. One might suppose that this is indifferent to the original, but this would lead to a contradiction. To see this, consider a slightly different way of shifting payoffs. To start, pick an arbitrary person with -1 payoff and shift the two adjacent 1s to her. The result is <…1,-1,1,-1,0,1,0,-1,1,-1,1,…>. Suppose, again, that this is indifferent to the original, since there was only a finite rearrangement of the payoffs. Now shift each -1 to the right of both 0s one position to the right, and shift each -1 to the left of both 0s one position to the left. The result is <…0,0,0,1,0,0,0,…>. Again, since the result of each individual shift was indifferent to the original, one might suppose that the result of an infinite number of shifts was indifferent to the original. This, however, leads to a contradiction: <…1,-1,1,-1,1,-1,…> would be indifferent to <…0,0,0,…> and to <…0,0,0,1,0,0,0…>, but the latter is better (by dominance) to the former.

Obviously, the issue is complex, but we hope that we have said enough to indicate that the strangeness of the result of an infinite number of indifferent shifts not being indifferent with the original is not peculiar to our particular approach. It is instead an inevitable result of some very weak assumptions in the context of aggregation over an infinite number of parts.

Let us here acknowledge what we take to be the weakest aspect of our approach. Although we are convinced that it is appropriate to appeal the above way to the natural order, if any, of *locations* of value (e.g., temporal location), it is far from obvious that it is appropriate to appeal to the natural order of the *constituents* of value (e.g., probability and payoff). We fully agree that sometimes it is not. For example, consider a version of weighted utilitarianism where the benefits of different individuals have different weights and the weights are stipulated to sum to one. Here the total value would be the sum of the weighted benefits, and the value for a given individual would have two constituents: her weight and her benefit. We fully agree that there is no reason to limit to admissible covering sequences to ones that are convex or symmetric in this case.

In the case of risky options, however, it seems that the structure of the constituents of value (probability and payoff) should be taken into account. We don't have a compelling argument for this, but the following consideration lends at least some support. Let $E_n(A)$ be the expected value of the risky option $A = \{<p_1,v_1>,<p_2,v_2>,\ldots,<p_n,v_n>,\ldots\}$ (enumerated in non-decreasing $|v_n|$) conditional on the *absolute value of the payoff* being no greater than $|v_n|$. For example, for the Pasadena game, P, $E_1(P)$ is 2 (= [½ x 2/1] x 2/1) and $E_2(P)$ is 8/3 (= [(1/2 x 2/1) + (1/4 x -4/2)] x4/3). One natural way of assessing the value of an option, A, is to take the limit, as n goes to infinity, of $E_n(A)$. This takes the limit of the conditional expected payoff as the cutoff on the absolute value of the payoffs is increased to infinity. A second natural way of assessing an option, $A = \{<q_1,w_1>,<q_2,w_2>,\ldots,<q_n,w_n>,\ldots\}$, now enumerated in non-increasing probabilities, is on the basis of the limit, as n goes to infinity, of $F_n(A)$, where $F_n(A)$ is the expected value of A conditional on the *probability* of the outcome being at least q_n. For example, for the Pasadena game, P, $F_1(P) = 2$ and $F_2(P) = 8/3$ (see E_1 and E_2 above) and $F_3(P)$ is 20/21 (= [(1/2 x 2/1)

+ (1/4 x -4/2) + (1/8 x 8/3] x 8/7). The limits of these two sequences of conditional expected values are the two most natural ways of assessing the value of options, when there is no absolute convergence of the summed values. Our approach agrees that these are admissible assessments, since, for a given option, each of the two expected value sequences corresponds to the value of the option relative to some admissible covering sequence. Our approach allows, however, that there can be other admissible covering sequences. It is thus weaker than appealing simply to the limits of these two sequences of conditional expected value. Obviously, this is not a compelling argument, but we hope that it at least helps motivate the approach.

We shall now compare this approach to the evaluation of conditionally convergent options with one recently developed by Kenny Easwaran.

3. Comparison with Easwaran's Weak Expected Value

Kenny Easwaran (2008) has tentatively suggested (in agreement with our approach) that the Pasadena game has a value of log2 on the following basis: (1) He distinguishes between weak expected value and strong expected value (defined below). (2) He establishes that the Pasadena game has a weak expected value of log2 even though its strong expected value is undefined (because of conditional convergence). (3) He tentatively suggests that in general risky options can be assessed on the basis of their weak expected value.

Easwaran is right that the Pasadena game has a value of log2. We shall argue, however, that, *in general*, risky options cannot be assessed merely on the basis of their weak expected value.

Easwaran appeals to two versions of the law of large numbers to define weak and strong expected value for a random variable. Let us consider a random variable, X, and let E(X) be its standard expected value. We shall say that E(X) exists just in case the sum of the probability-weighted payoffs is well defined and equal to a finite real number. (This is true just in case the expected value of the absolute value of X, E(|X|), is finite.). E(X), for example, does not exist in the Pasadena game. Let Ave(X,n) be the average value of X for n repeated independent trials. The two laws are:

Strong Law of Large Numbers: If E(X) exists, then, for each positive number e, prob(the limit as n goes to infinity of |Ave(X,n) – E(X)| < e) = 1.

Weak Law of Large Numbers: If E(X) exists, then, for each positive number e, the limit as n goes to infinity of prob(|Ave(X,n) – E(X)| < e) = 1.

Both laws concern the probability of the expected value and the average value being arbitrarily close to each other. The difference between the two laws concerns whether the limit as the sample size goes to infinity is internal to the probability assignment or external to it. Easwaran uses these two laws to define the weak and strong expected value as follows:

Strong Expected Value: The strong expected value of a random variable, X, is the number, v, if any, for which for any positive number e, prob(the limit as n goes to infinity of |Ave(X,n) – v| < e) = 1.

Weak Expected Value: The weak expected value of a random variable, X, is the number, v, if any, for which for any positive e, the limit as n goes to infinity of prob(|Ave(X,n) – v| < e) = 1.[9]

The strong expected value is the value v for which there is *a probability of 1 that, for large enough sample sizes,* the average value will be arbitrarily close to v. The weak expected value, by contrast, is the value v for which, *for large enough sample sizes, there is a probability of 1* that the average value will be arbitrarily close to v. Whenever the strong expected value is defined, then the weak expected value is also defined and has the same value. The weak value, however, can exist when the strong value does not.[10] Indeed, Easwaran proves that, in the Pasadena game, the weak expected value is log2, but the strong expected value is not defined.

Easwaran tentatively proposes the following principle:

Weak Expectations: The value of an option is its weak expected value.

Although we agree with Easwaran that the value of the Pasadena game is log2, and we agree that the appeal to weak expected value, when it exists, is relevant, we shall argue that: (1) Weak expected value does not fully determine the value of an option. (2) Options can be evaluated even when the weak expected value does not exist. In short, appealing to the weak expected value is not, as Easwaran suggests, central to evaluating conditionally convergent risky options.

We fully agree that the weak expected value, when it exists, is relevant to assessing risky options. This is because, when it is exists, the weak expected value is equal to the limit of a certain admissible (i.e., convex and quasi-symmetric) covering sequence. As Easawarn notes, the weak expected value of an option A, when it exists, is the limit of the sequence $\langle E(A_1), E(A_2), \ldots E(A_n), \ldots\rangle$, where $E(A_n)$ is the expected value of an option exactly like A, except that all payoffs with an absolute value greater than n are reduced to 0. The sequence thus starts by ignoring payoffs with absolute values greater than 1, and then progressively takes into account payoffs for which the absolute value is higher. At the limit, it takes into account all payoffs. This, of course, is equivalent to the limit of the values of A relative to the convex quasi-symmetric covering sequence $\langle S_1, S_2, \ldots, S_n, \ldots\rangle$, where S_n is the set of $\langle p_n, v_n\rangle$ for which $|v_n| \leq n$. Hence, we fully agree that the weak expected value, when it is exists, is relevant to the assessment of risky options.

Easwaran supports his tentative endorsement of Weak Expectations by noting that a player who plays a game a very large number of times at a price that is slightly higher (respectively: lower) than the weak expectation has a very high probability of ending up behind (respectively: ahead). Indeed, by repeating the game enough times, that probability can be made as close to 100% as one likes.[11]

We agree that the weak expectation, when it exists, is one of the relevant measures of an option. We deny, however, that it is the only one. Thus, if (as in the Pasadena game) the weak expected value is finite, and one pays less than that for each play of the game, by repeating the game enough, one can make the probability arbitrarily high that one will be ahead *at any given stage* after some point. Nonetheless, the chance of being ahead *at all (infinitely many)* further stages is zero (Durrett 2005, Ch. 1, Theorem 6.7). That is, although a finite weak expected value, wev, is such that, for any positive ε and δ, there is a k, such that for all n≥k, prob(|ave(n) – wev| < ε) ≥ 1-d, where the standard expected value does not exist, it is nonetheless true that, for the same ε, δ, and k, prob(for all n≥k, |ave(n) – wev| < ε) = 0. This is because the latter requires that the average be close enough to the weak expected value for an infinite number stages of the game, and, where the standard expected value does not exist, this probability is 0. In other words, in these cases, for any arbitrarily low negative number, the probability is 1 that there are infinitely many stages of play for which the average net gains from play are below that negative number.

Although this does not provide a reason to reject completely the relevance of weak expectations, it does provide a reason to deny that the value of an option is fully determined by its weak expectation. For example, the value of the modified (probability-splitting) Pasadena option introduced in the previous section has a weak expectation of log2. Although that is an admissible evaluation, it is not the only admissible evaluation. Other admissible covering sequences lead to other values. Because no admissible convex quasi-symmetric covering sequence is privileged, this option, we claim, has indeterminate value. It is at least as valuable as .5xlog2 and not more valuable than log2, but there may be no more determinate fact about its value.

To put this in perspective, we note that our approach agrees that the value of an option is determinately its weak expected value in the special case where the option is such that (1) the weak expected value exists and is finite, and (2) the orderings of the probability-payoff pairs, $<p_n,v_n>$, according to increasing absolute payoff is the same, perhaps with finitely many exceptions, as the ordering according to decreasing probabilities. This is because the existence of finite weak expected value entails that the product xProb(|X|>x) goes to 0 as x goes to infinity (Durrett 2005, Ch. 1, Remark to 5.6), and that entails Thin Tails. Thus, these two conditions ensure that the conditions of Limit Agreement 2 (above) hold and thus that the value of the option is the same for admissible covering sequences. Given that the weak expected value is the value of one of those sequences, it is the unique value. Thus, we fully agree with Easwaran in this special case.

Above we argued that Weak Expectations is too strong in a certain respect. We now argue that it is too weak in another. Weak Expectations evaluates an option only when it has a weak expected value. We see, however, no reason to restrict evaluations to such cases. Consider, for example, an option consisting of all the <probability, payoff> pairs of the form, for each natural number k, $<1/Kxk^2, (-1)^{k+1}Kxk>$, where K is the sum of $1/k^2$ (= $\pi^2/6$ or about 1.645). This option generates the same conditionally convergent series, 1 -1/2 +1/3 …, as the Pasadena game. Furthermore, as with the original Pasadena game, for each admissible covering sequence, the limit of the values of this option is log2. Thus, it seems perfectly appropriate to hold that log2 is its value. The option, however, does not satisfy condition that xProb(|X|>x) goes to 0 as x goes to infinity. This condition does not hold for this option, since $vP(|V|>v) = (Kxk)x(p_{k+1}+p_{k+2}+\ldots) > (Kxk)x(p_{k+1}+p_{k+2}+\ldots+p_{2k}) > (Kxk)x(k/(Kx(2k)^2) = 1/2^2 = .25$.

Thus $vP(|V|>v)$ does not go to zero as v goes to infinity. Thus, we see no reason to restrict the assessment of conditionally convergent options to cases where the weak expected value exists.

Thus, although we agree with Easwaran that the value of the Pasadena game is log2 and the weak expectation is indeed relevant to the assessment of options when it exists, we deny that (1) the weak expectation, when it exists, always fully determines the value of an option, and (2) that no evaluation of options is possible when the weak expectation does not exist.[12]

4. Extensions

Before closing, we shall identify, without defense, three ways in which our proposed principle, Weak Generalized Catching-Up, can be plausibly strengthened.

4.1 Applying the Principle Where the Sums Are Not Conditionally Convergent

We have focused throughout on conditionally convergent sums, but the principles developed apply equally well for other kinds of sums that fail to converge absolutely. A sum fails to converge absolutely just in case the order of summing affects the result. Conditional convergence is the special case where the terms that are summed converge to zero (e.g., 1, -1/2, 1/3, …). Let us therefore briefly address the cases where they do not converge to zero (e.g., 2, -1, 2, -1, …).

Consider, for example, a game consisting of the <probability, payoff> pairs of $<p_k, 2/p_k>$ for odd k, and $<p_k, -1/p_k>$ for even k, where p_k is, as in Pasadena, $1/2^k$. In increasing order of the absolute value of payoffs, this gives probability-weighted payoffs of …-1,-1,-1, 2, 2, 2, … No matter what convex set of locations one starts with, all convex quasi-symmetric expansions have a limit of positive infinity. (For example, no matter how far to the “left” one starts in this sequence, the total value relative to convex quasi-symmetric expansions will eventually become positive and then increase without bound.) Hence, we claim that such a gamble is more valuable than any gamble with a finite expected payoff.

It's worth noting here that if the positive payoffs above were <p_k, $1/p_k$> rather than <p_k, $2/p_k$>, then the assessment would be radically indeterminate. In increasing order of the absolute value of payoffs, this gives probability-weighted payoffs of …-1, -1, -1, 1, 1, 1, …, and the total value relative to convex expansions can be made as negative (respectively: positive) as one wants by starting far enough to the left (right).

4.2 Admissible Starting Points for Symmetric Distributions

For Weak Generalized Catching-Up, there are no privileged starting points. All possible starting points must be considered. A second plausible strengthening is to privilege "symmetric starting points" in the special case where the distributions are "symmetric". Let us say that a probability distribution over payoffs is *symmetric in payoffs* just in case there is some payoff, v, such that, for any t, the probability of payoff v+t is the same as the probability of payoff v-t. For distributions that are symmetric in payoffs, there cannot be more than one point, v, of symmetry (since otherwise probabilities would sum to more than one). Thus, we may refer to *the point of symmetry*. (Probability distributions over infinitely many payoffs can't be *symmetric in probabilities*, and so we ignore that case.)

Consider the following symmetric distribution, which has 0 as its point of symmetry:

Positive payoffs pairs: <¼, 4>, <1/8, 8>, <1/16, 16>, ….

Negative payoff pairs: <¼, -4>, <1/8, -8>, <1/16, -16>, ….

In increasing order of absolute payoff, this generates the probability-weighted values of <…-1,-1,-1,1,1,1…>. If no starting point is privileged, then, by starting sufficiently to the left (right) of the point of symmetry, the limit of the partial sums can be made as negative (positive) as one likes. This suggests that this option has completely indeterminate value. We claim, however, that it has value 0, since it has zero as its unique point of symmetry. More generally, we believe that an option that has v as its unique point of symmetry in payoffs has value v.[13] More generally still, we believe that, for options that are symmetric in payoffs, only symmetric starting points are relevant.

We thus propose that only admissible starting points are relevant, where a starting point is admissible for an option only if (1) the point is a point of symmetry, and the option is symmetric in payoffs, or (2) it is any point, if the option is not symmetric in payoffs. This will give stronger results, which, we believe, are plausible.

4.3 Extending the Principle to Differentiate Among Infinitely Valuable Options

Consider an option for which there is, for each natural number n, a $1/2^n$ probability of winning 2^n and no probability of losing anything or winning anything else (the St. Petersburg gamble). All admissible covering sequences will assign an infinite value to this option, and hence its value is infinite. Consider now a second option, just like the first, except that for each $1/2^n$ chance one wins 2^n+1 (rather than 2^n). All admissible covering sequences assign this an infinite value as well. As a result, Weak Generalized Catching-Up is silent, since (1) it judges one option as better than another only if for at least one admissible covering sequence the former has a higher limit than the latter (but both are infinite), and (2) it judges two options equally valuable only when each admissible covering sequences assign the same *finite value* to each option. Still, it seems clear that the second option, which dominates the first, is more valuable. Hence, the following strengthening seems plausible.

Weak Generalized Catching-Up cannot distinguish between two options both of which have infinite limits relative to all the admissible covering sequences. When one option dominates the other, this, we claim, is inappropriate. The revision below replaces (1) the appeal to the comparison of *the limits* of the values of two options relative to a given admissible covering sequence with (2) an appeal to whether the value of one option is at least as great, and perhaps greater, than the value of the other option, *relative to all but finitely many of the sets* in the covering sequence ($<S_1, S_2, S_3, \ldots>$). For example, in the above example, for any admissible covering sequence, and any set in that sequence, the value of the second option (with the 1 unit higher payoff) is higher than the first option.

Consider, then:

Generalized Catching-Up: For any two boundedly finite options, A and B:
(1) A is *more valuable* than B if and only if (a) for *each* admissible covering sequence, $<S_1,S_2,\ldots,S_k,\ldots>$, for *all but finitely many* S_i, the value of A relative to S_i is *greater than or equal* to the value of B relative to S_i, and (b) for *some* admissible covering sequence, $<R_1,R_2,\ldots,R_k,\ldots>$, there is a positive number, e, such that for *infinitely many* R_i, the value of A relative to R_i is *greater than* e plus the value of B relative to R_i;
(2) A is *equally valuable* with B if and only if for *each* admissible covering sequence, $<S_1,S_2,\ldots,S_k,\ldots>$, and each positive number e, for *all but finitely many* S_i, the absolute value of difference between the value A relative to S_i and the value of B relative to S_i is less than e.[14]

To illustrate the content of these principles, let us suppose that there are only two admissible covering sequences and the values of the options, relative to these sequences, are as follows:

	Value Relative to Covering Sequence 1	Value Relative to Covering Sequence 2
A	<1,2,3,4,5,6,…>	<1,2,3,4,5,6,…>
B	<2,2,3,4,5,6,…>	<1,1,3,3,5,5,…>

Relative to the sets in these sequences, the value of A is always at least as great as that of B, except in a finite number of cases: the first set of the first sequence (1 vs. 2). The fact that the value of A is *lower* than that of B relative a *finite* number of sets, in a given covering sequence, does not cast doubt on A being more valuable than B. After all, after some finite stage in the covering sequence (here: the first stage), relative to any subsequent set, the value of A is always at least as great as that of B. Likewise, the fact that the value of A is *higher* than that of B relative to a *finite* number of sets, for each covering sequence, is not sufficient to establish that A is more valuable than B. If, however, relative to each admissible covering sequence, the value of A is lower than that of B in only a *finite many* sets and greater in *infinitely many* sets, then A is more valuable than B. In Sequence 2, A is more valuable than B relative to infinitely many sets. Hence, A is more valuable than B.

The above principle is a strengthening of the original in three ways. First, we have dropped the appeal to the limits, and hence the principle has a wider domain of applicability. For example, suppose that, relative to a given covering sequence, the values of A are <-1,1,-1,1,…> and the values of B are <-1,-1,-1,…>. The second sequence has a limit of -1, but the first sequence has no limit (instead it has two cluster points, -1 and 1). Nonetheless, the extended principle can assess these two options relative to this sequence, since, for all sets, the value of A is at least as great as that of B, and, for infinitely many sets, it is greater. Second, where A and B are each infinitely valuable, the original principle was silent, whereas the strengthened versions sometimes will judge one better than another. The example of the previous paragraph illustrates this. The example with which this section started also illustrates this. Here, A has a payoff of 2^n with probability $1/2^n$ and B has payoff of 2^n+1 with probability $1/2^n$. Both are infinitely valuable and thus the Weak Generalized Catching-Up was silent. The strengthened version, however, judges B more valuable, since, for every set of every admissible covering sequence, the value of B restricted to this set is more than .5 greater than the value of A. Third, the revised principles provide necessary and sufficient conditions for one option being at least as valuable as another, whereas the original version only provided sufficient conditions. Generalized Catching-Up does not generate a complete ordering, but it is well known in the literature that completeness in an infinite world is hard to obtain (e.g. Lauwers, 2010). Our claim is that no stronger principle is plausible.

It's important to note here that, because the principle has been strengthened to provide necessary, as well as sufficient, conditions for judging one option at least as good as another, when applying this principle, covering sequences must satisfy *all relevant conditions of admissibility.* We have claimed that these conditions include (1) convexity, (2) quasi-symmetry, and (3) centeredness on the payoff point of symmetry if the option is symmetric with respect to payoffs, but we have not claimed that they exhaust those conditions. Our working assumption is that they do, but this is a matter for further investigation.

It's worth noting that Generalized Catching-Up can be restated more simply by appealing to the technical notions of cluster point and the limit inferior of a sequence. We allow cluster points to be finite or infinite. A *finite cluster point* of a sequence is a value for which there are *infinitely* many values in the sequence that are arbitrarily close to it. A sequence has a cluster point of positive (respectively: negative) *infinity* just in case, for any positive (negative) real number n, there are *infinitely* many values in the sequence that are greater (less) than n. The *limit inferior* of a sequence is the smallest cluster point. A sequence has a *limit* just in case there is exactly one cluster point. For example, <1/2, ¼, ¾, 1/8, 7/8, 1/16, 15/16, …> has two cluster points, 0 and 1, and hence no limit. Its limit inferior is 0.

It's straightforward to establish that Generalized Catching-UP can restated as follows:

Generalized Catching-Up: A is *at least as valuable as* B if and only if, for *each* admissible covering sequence, $<S_1,S_2,\ldots,S_k,\ldots>$, the limit inferior of [the value of A minus the value of B], relative to the sets in this sequence, is *greater than or equal to 0.*

To illustrate this, reconsider the above example:

	Value Relative to Covering Sequence 1	Value Relative to Covering Sequence 2
A	<1,2,3,4,5,6,…>	<1,2,3,4,5,6,…>
B	<2,2,3,4,5,6…>	<1,1,3,3,5,5,…>
A-B	<-1,0,0,0,0,0,…>	<0,1,0,1,0,1,…>
B-A	<1,0,0,0,0,0,…>	<0,-1,0,-1,0,-1,…>

Relative the first covering sequence, the value of A minus the value of B has a limit of 0 (since all but one of the differences in value is 0). Relative to the second covering sequence, the value of A minus the value of B has two cluster points, 0 and 1. For each sequence, the limit inferior is thus 0. Thus, A is at least as valuable as B. B, however, is not at least as valuable as A, since, relative to the second sequence, the limit inferior of the value of B minus the value of A is -1 (and hence less than 0). Thus, A is more valuable than B (in accordance with the judgment of the original formulation).

5. Conclusion

Where, as in the Pasadena game, the standard expected value of an option is not absolutely convergent, standard value theory (e.g., decision theory) is silent about the evaluation of that option. We suggested that where the payoffs are ordered over time (which has a natural order), an evaluation on the basis of the temporal order may be appropriate. This, however, is a very special case and it is silent about the standard Pasadena game, which need not involve any temporal dimension. We suggested that options should be evaluated on the basis of agreement among admissible (e.g., convex and quasi-symmetric) covering sequences of the constituents of value (i.e., probabilities and payoffs). It's crucial to keep in mind that the proposed principles are advocated as normative (e.g., prudential) principles and not as mathematical principles. The mathematical sums are clearly undefined in cases of conditional convergence.

In the case of the Pasadena game, our approach agrees with Easwaran's weak expectations approach that the value is log2. We have argued, however, that it is possible to evaluate options even when there is no weak expectation, and that, even where it exists, the weak expectation does not, pace Easwaran, fully determine the value of an option (although it is relevant).

Obviously, these claims go beyond standard value (e.g., decision) theory. They are substantive normative claims about the evaluation of options. Although we have motivated these claims, we haven't given conclusive reason to accept them. We hope, however, that we've said enough for them to be taken seriously and investigated further.[15]

References

Alexander, J. McKenzie (2010). "Decision theory meets the Witch of Agnesi," unpublished draft.

Atsumi, Hiroshi (1965)." Neoclassical Growth and the Efficient Program of Capital Accumulation," *The Review of Economic Studies*, 32: 127-136.

Baker, Alan (2007). "Putting Expectations in Order," *Philosophy of Science* 74: 692–700.

Colyvan, Mark (2008). "Relative Expectation Theory," *Journal of Philosophy*, 105: 37–44.

Easwaran, Kenny (2008). "Strong and Weak Expectations," *Mind* 117(467):633-641

Durrett, Rick (2005). *Probability: Theory and Examples*, 3rd edition (Belmont, CA: Thomson Brooks/Cole).

Feller, William (1971). *An Introduction to Probability Theory and Its Applications,* volume 2 (New York: John Wiley & Sons Inc.).

Terrence L. Fine, "Evaluating the Pasadena, Altadena, and St. Petersburg Gambles," *Mind* 117 (2008): 613-32.

Fleurbaey, Marc and Philippe Michel (2003). "Intertemporal Equity and the Extension of the Ramsey Criterion," *Journal of Mathematical Economics* 39(7), 777-802.

Hájek, Alan (2009). "All Values Great and Small," manuscript.

Hájek, Alan and Harris Nover (2006). "Perplexing Expectations," *Mind* 115 (July), 703-720.

Hájek, Alan and Harris Nover (2008). "Complex Expectations," *Mind* 117 (July), 643-664.

Lauwers, Luc (2010). "Ordering infinite utility streams comes at the cost of a non-Ramsey set," *Journal of Mathematical Economics* 46, 32-37.

Lauwers, Luc and Peter Vallentyne (2004). "Infinite Utilitarianism: More Is Always Better," *Economics and Philosophy* 20: 307-330.

Nover, Harris and Alan Hájek (2004). "Vexing Expectations," *Mind* 113 (April), 237-249.

Smith, Nicholas (2010). “Is Evaluative Compositionality a Requirement of Rationality?” unpublished draft.

Sprenger, Jan and Remco Heesen (2009). “The Bounded Strength of Weak Expectations,” unpublished.

Vallentyne, Peter and Shelly Kagan (1997). “Infinite Utility and Finitely Additive Value Theory,” *Journal of Philosophy* 94: 5-26.

Carl Christian Von Weizsäcker (1965). “Existence of Optimal Programs of Accumulation for an Infinite Time Horizon,” *Review of Economic Studies* 32: 85-104.

Appendix

We here outline the proofs (assuming locations are discrete) for the limit agreement results in the paper:

1. Proof of Limit Agreement 1

Limit Agreement 1: All quasi-symmetric convex covering sequences of a given option have the same limit when all the following conditions holds: (1) the basic locations have a natural order and a natural distance metric, (2) there is only one dimension for basic locations (e.g., just time), (3) the option is boundedly finite, and (4) for any location, the values in other locations converge to zero as their distance from the given location tends to infinity.

Proof: Consider an option $O = \langle \ldots, v_{-n}, \ldots, v_{-2}, v_{-1}, v_0, v_1, v_2, \ldots, v_n, \ldots \rangle$, where v_k is the value at location k. Assume that v_k and v_{-k} both converge to zero as k goes to infinity. Consider the covering sequence of locations, $\langle S_0, S_1, S_{2, \ldots} \rangle$ for which $S_0 = \{0\}$, $S_1 = \{-1,0,1\}$, $S_2 = \{-2,-1,0,1,2\}$, …. The sequence of value for O, relative to this quasi-symmetric sequence, $\langle O(S_0), O(S_1), O(S_2), \ldots \rangle$ is $\langle V_0, V_1, V_2, \ldots \rangle$, where $V_k = v_{-k} + \ldots + v_{-2} + v_{-1} + v_0 + v_1 + v_2 + \ldots + v_k$. Suppose that this sequence converges to the value v (finite or infinite) as k goes to infinity. The summations are symmetric around t=0. Now, shift the center in the covering sequence from 0 to n with n > 0. Obtain the new covering sequence $\langle \{n\}, \{n-1, n, n+1\}, \{n-2, n-1, n, n+1, n+2\}, \ldots \rangle$ and the corresponding sequences of value for O relative to this sequence, $\langle W_0, W_1, W_2, \ldots \rangle$, where $W_k = v_{n-k} + \ldots + v_{n-2} + v_{n-1} + v_n + v_{n+1} + v_{n+2} + \ldots + v_{n+k}$. The difference $|W_k - V_k|$ is less than 2n x maximum$\{|v_{-k}|, \ldots, |v_{n-k+1}|, |v_{k+1}|, \ldots, |v_{n+k}|\}$. This maximum goes to zero as k goes to infinity (recall that v_k and v_{-k} converge to zero as k goes to infinity). Hence, $|W_k - V_k|$ converges to zero as k goes to infinity. Thus, W_k converges to v as well. The summation, that is, does not depend upon the starting point (t=0, versus t=n). The above argument considers minimal expansions (each set in the covering sequence originates from its predecessor by adding two moments in time). Obviously, the result extends to arbitrary convex and quasi-symmetric expanding sequences.

2. Proof of Limit Agreement 2

Thin Tails: For each positive real number r, the limit of the sum of probability-weighted absolute values, $p_k|v_k|$, for $x \leq |v_k| \leq x+r$, goes to zero as x goes to infinity.

Limit Agreement 2: Let $O = \{<p_1,v_1>,<p_2,v_2>,\ldots,<p_n,v_n>,\ldots\}$ be a risky option, enumerated on the basis of non-decreasing *absolute values* ($|v_n| \leq |v_{n+1}|$, for each n). Suppose that the ordering of these pairs according to non-increasing *probabilities* is the same, perhaps with finitely many exceptions, as the ordering listed above, and that $p_1v_1+ p_2v_2+\ldots+ p_nv_n+\ldots$ converges to V (finite or infinite). If Thin Tails is satisfied, then the limits of the sum of the probability-weighted values of this option are the same for all convex and quasi-symmetric covering sequences. In addition, if V is a finite number, then all these sums are the same *if and only if* condition Thin Tails holds.

Proof: We proceed in four steps.

Step 1: Set-up and notation

Since the orderings of the pairs, $<p_n,v_n>$, according to non-decreasing $|v_n|$ and according to non-increasing p_n are the same up to finitely many exceptions, there exists a natural number N such that from N onwards the two orderings coincide exactly and such that p_N is below (or equal to) the probabilities $p_1,p_2,\ldots,p_{N-1}$.

Throughout we consider convex and quasi-symmetric covering sequences. In order to study the limiting behavior of the value of the option, restricted to the sets that belong to such a covering sequence, we assume (without loss of generality) that the first set in such a covering sequence already contains the points $<p_1,v_1>,<p_2,v_2>,\ldots,$ and $<p_N,v_N>$.

Step 2: The value relative to an arbitrary covering sequence

Let $S_1, S_2, \ldots, S_k, \ldots$ be a covering sequence. Quasi-symmetry implies that the sets S_k are symmetric with respect to some payoff-line, v. Assume that $v > 0$ (the cases $v < 0$ or $v=0$ are analogues). We show that, according to this sequence, the value of the option O is the limit of O[2v-x,x] as x goes to infinity, where O[r,s] is a shorthand for the sum of the probability weighted values of the $<p,v>$ pairs that belong to the option O and that have value v between r and s (i.e., $r < v < s$).

Consider a set $S=S_t$ in the covering sequence. Let $<p_j,v_j>$ with $j \geq N$, be the pair with the *highest* payoff that belongs to the intersection of the option O and the set S. Due to the symmetry with respect to the payoff-v-line, the pair $<p_j ,v-(v_j-v)> = <p_j ,2v-v_j>$ also belongs to S. Quasi-symmetry further requires that S is symmetric with respect to the .5 probability-line. Thus the pairs $<1-p_j ,v_j>$ and $<1-p_j ,2v-v_j>$ belong to S. Here, assume that v_j is large enough for $2v-v_j$ to be negative (otherwise, consider a superset of S further in the covering sequence). Due to convexity, the rectangle spanned by the four points $<p_j,v_j>$, $<p_j , 2v-v_j>$, $<1-p_j , v_j>$, and $<1-p_j , 2v-v_j>$ is a subset of S. Because the orderings in decreasing probabilities and in increasing absolute payoffs coincide, and $j \geq N$, all points $<p_i,v_i>$ in O with positive payoff smaller than v_j (and probability larger than pj) are included in S.

Similarly, let $<p_m , v_m>$ with $m \geq N$, be the pair with the *lowest* (negative) payoff that belongs to O and to S. Then, all points $<p_i,v_i>$ in O with negative payoff larger than v_m (and probability larger than p_m) are included in S. Hence, the value of the option O restricted to S is equal to O[v_m,v_k]. Let y>0 be the maximum of $|v_m – v|$ and $|v_k – v|$. Then, $O[v_m,v_k]$ = O[v-y , v+y] = O[2v-x, x] with x equal to v+y.

Step 3: The agreement of each covering sequence

Consider a covering sequence. Let the value-v-line be its line of symmetry (with v>0). From the previous step, we know that this sequence assigns the limit of O[2v-x,x] as x goes to infinity as the value of the option O.

From the definition of V, we see that V is the limit of O[-x,x] as x goes to infinity. The difference between O[2v-x,x] and O[-x,x] is captured by O[-x,2v-x]. If Thin Tails is satisfied, then this difference O[-x,2v-x] goes to zero as x goes to infinity. Thus, each covering sequence proposes the same value.

Step 4 The necessity of Thin Tails for finite V

Consider a covering sequence with v-payoff line as its line of symmetry, with v>0. This covering sequence assigns the limit of O[2v-x,x] as x goes to infinity as the value of the option O. From the previous step, we know that this is equal to the limit of O[-x,x], which is V. Since, V is finite the limit of the difference O[-x,2v-x] must converge to zero. Likewise, if the line of symmetry is the v-payoff line for v≤0, then O[x-2v,x] must converge to zero. Either way, Thin Tails is satisfied.

[1] An example of an ordered sum that does not have a limit (finite or infinite) is 1-1+1-1+1… It oscillates between 1 and 0.

[2] Not all ordered sums with infinite limits are order-dependent. They are so if and only if (1) they involve terms of the opposite sign as the limit, and (2) the sum of those terms, in their specified order, does not converge to a finite number. For example, the ordered sum of 1,1,1,1,1… and the ordered sum of 1,-1/2,1,-1/4, 1,-1/8,… each converge to positive infinity but are not order-dependent.

[3] We thus deny that the standard axioms of rational choice exhaust the relevant axioms. For discussion of this issue in the context of the Pasadena game, see Fine (2008).

[4] These examples also show a standard continuity condition (on the product topology) is violated: For any A, if O is the limit of O_1, O_2, O_3, …, and each $O_i \geq A$, then $O \geq A$. Let A be <1,0,0,0,0,…>. Each O_i is better than A and thus continuity requires that the limit of the sequence, <0,0,0,…> be at least as good as A, but it is worse (because dominated). Note that these are not examples of intransitivities.

[5] See also Hájek and Nover (2006), Hájek and Nover (2008), and Hájek (2009).

[6] See Baker 2007 for appeal to the temporal order in which payoffs are made.

[7] More generally, for n-dimensional space, a set is m-dimensionally symmetric (for $0 \leq m \leq n$) just in case the maximal number, up to n, of n-1dimensional hyperplanes relative to which the set is symmetric is m. (An n-dimensional hyperplane is a point for n=0, a line for n=1, and a plane for n=2.)

[8] More generally, for n-dimensional space, where the set of locations is m-dimensionally symmetric, a covering sequence is *quasi-symmetric* just in case, if $m \geq 1$, there are m intersecting n-1dimensional hyperplanes of symmetry for the space of locations relative to each of which each set in the sequence is symmetric.

[9] These are definitions are definitions of *finite* strong/weak expected value. The notions can be extended to include *infinite* strong/weak expected value as follows (with the definitions for negative infinity left implicit): (1) Strong Expected Value (positive infinity): The strong expected value of a random variable, X, is positive infinity just in case, for each positive number e, prob(the limit as n goes to infinity of Ave(X,n) > e) = 1; (2) Weak Expected Value: The weak expected value of a random variable, X, is positive infinity just in case, for each positive e, the limit as n goes to infinity of prob(Ave(X,n) > e) = 1. For simplicity, however, in the text, we shall follow Easwaran and assume that the values are finite.

[10] As a matter of fact the following holds. Let X be a random variable and let $X_n = X\, 1_{(|X|\leq n)}$ denote the truncated variable, i.e. X_n coincides with X if the absolute value is less than or equal to n and is set equal to 0 otherwise. The weak expected value of a random variable X exists and is finite *if and only if* (i) the product xProb(|X|>x) converges to 0 as x goes to infinity, and (ii) the sequence $E[X_1]$, $E[X_2]$, …, $E[X_n]$, … of truncated expected values converges to a finite number as n goes to infinity. In addition, the weak expected value is equal to this limit of truncated expected values and, hence, is equal to the sum of the probability-weighted payoffs ordered on the basis of increasing absolute value of payoffs.

[11] For example, if the Pasadena game is repeated 250,000 times, then there is a 95% probability that the average will between log2 – 0.01 and log2 + 0.01. By repeating more times, the probability can be made as close to one as one likes.

[12] For additional discussion of Easwaran's approach, see Sprenger and Heesen (2009). They consider an approach that assumes that utility is bounded. For criticism and an alternative view, see Smith (2010).

[13] Here we have been influenced by a similar suggestion in Alexander (2010).

[14] Generalized Catching-Up is similar to Strengthened Basic Idea 3 of Vallentyne and Kagan (1997). It is stronger in that it appeals to convex sets rather than the weaker notion of bounded regions (the interior

of a closed curve, plus perhaps the curve). It is weaker in that it merely requires each expansion be quasi-symmetric rather than the more restrictive requirement that it expand “uniformly” in all directions from the center (which requires that the same shape be preserved).

[15] For extremely helpful comments, we thank [??? Removed for .