Symmetries of Value

Zachary Goodsell

March 13, 2025

Abstract

Standard decision theory ranks risky prospects by their expected utility. This ranking does not change if the values of all possible outcomes are uniformly shifted or dilated. Similarly, if the values of the outcomes are negated, the ranking of prospects by their expected utility is reversed. In settings with unbounded levels of utility, the expected utility of prospects is not always defined, but it is still natural to accept the *affine symmetry principles*, which say that the true ranking of prospects is unchanged by shifts and dilations, and is reversed by negation—even in hard cases where expected utilities are undefined. This paper investigates the affine symmetry principles and their consequences. The principles are found to be surprisingly powerful. Combined with orthodox axioms, they assign precise utility values to previously problematic cases: for example, $\ln 2$ to the Pasadena prospect (Nover and Hájek, 2004) and $-1/2$ to the alternating St Petersburg prospect. They also have important structural consequences, notably vindicating Colyvan's (2008) Relative Expectation Theory. Finally, the paper establishes the consistency of the affine symmetry principles. In light of their fruitful consequences, this result supports their adoption as fundamental axioms of decision theory.

1 The Affine Symmetry Principles

This paper investigates two natural principles of decision theory, together called the *affine symmetry principles*, which posit how certain symmetries in the value of the outcomes of risky prospects induce corresponding symmetries in the comparison of the prospects themselves. The principles, despite initially seeming obvious, are shown to have striking consequences for decision theory in cases where the standard expected utility of a prospect is undefined.

The possible outcomes of prospects have an interval scale structure, corresponding to how much a chance of them contributes to the overall value of a prospect. For example, we say that outcome o_2 is halfway between o_1 and o_3 if getting o_2 for sure is equally good as a fifty-fifty gamble between o_1 and o_3, whereas o_2 is a third of the way from o_1 to o_3 if getting o_2 for sure is equally good as a 2/3 chance of o_1 and a 1/3 chance of o_3. This structure can be graphically represented by placing the outcomes on a line as in Figure 1, so that the ordering of the outcomes is represented by the left-to-right ordering of points on the line, and the ratios of differences between outcomes are represented by the ratios of distances between their corresponding points.

Typical approaches to decision theory—including the widely accepted Expected Utility Theory—constrain the overall ranking of prospects in terms of the interval scale structure of outcomes. For any such approach, the comparison of two prospects depends only on the interval scale structure exhibited by the outcomes those prospects might yield. If Figure 2 is accurate about the interval scale structure of outcomes, this means that any verdict these approaches yield about gambles

Figure 1: Interval scale structure of outcomes represented by placing outcomes called "Death", "Coal", "Cake", and "Heaven" on the line. The diagram represents Heaven as the best, Death as the worst, certainty of Coal as being equally good as 3/4 chance of Death and a 1/4 chance of Heaven, certainty of Cake as being equally good as 2/3 chance of Coal and 1/3 chance of Heaven, and so on.

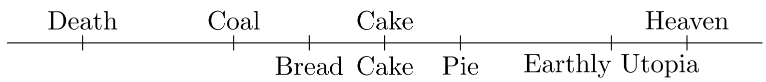

Figure 2: In this figure with three additional outcomes ("Bread", "Pie", "Earthly Utopia") are identified on the line, so that the outcomes with labels above the line exhibit the same interval scale structure as the outcomes with labels below the line. Both above and below the line, the second worst outcome is a quarter of the way from the worst to the best, the third worst is half of the way from the worst to the best, and so on.

between the outcomes Death, Coal, Cake, and Heaven would be the same if Death were replaced with Bread, Coal with Cake, Cake with Pie, and Heaven with Earthly Utopia.

A very natural idea, which is the topic of this paper, is that the comparison between any two prospects does not depend on any feature of outcomes besides their interval scale structure. This is the principle of *Positive Affine Invariance*. The other idea investigated in this paper is closely related. The principle of *Negative Affine Anti-Invariance* says that if the ordering of prospects by *better than* and the ordering of prospects by *worse than* are determined by the interval scale structure of the outcomes in the same way, which is to say that reversing the interval scale structure while keeping all ratios of differences the same should flip the comparison of any two prospects. Together these are called the *affine symmetry principles*.

To state the affine symmetry principles more carefully, let an *affine permutation* of outcomes be a permutation of outcomes π that preserves the ratios of differences between outcomes: so that $\pi(o_2)$ is the same proportion of the way from $\pi(o_1)$ to $\pi(o_3)$ as o_2 is from o_1 to o_3. An affine permutation is said to be *positive* if it preserves the ordering of outcomes: $\pi(o_1)$ is at least as good as $\pi(o_2)$ just in case o_1 is at least as good as o_2; otherwise it *reverses* the ordering of outcomes, in which case the affine permutation is said to be *negative*.

For an arbitrary prospect μ let $[\pi]\mu$ be the prospect that yields $\pi(o)$ whenever μ would yield o. The affine symmetry principles may now be stated as follows:

Positive Affine Invariance For any positive affine permutation of outcomes π, and for any prospects μ and ν, $\mu \succsim \nu$ if and only if $[\pi]\mu \succsim [\pi]\nu$.

Negative Affine Anti-Invariance For any negative affine permutation of outcomes π, and for any prospects μ and ν, $\mu \succsim \nu$ if and only if $[\pi]\mu \precsim [\pi]\nu$.

Which is to say that the comparison of prospects μ and ν is insensitive to anything besides the interval scale structure of the outcomes they might yield.

The standard view in decision theory, Expected Utility Theory, includes theorems very close to the affine symmetry principles. A *utility representation* is a function from outcomes to real numbers

Table 1: The alternating St Petersburg prospect

$1/2$	$1/4$	$\dots$	$1/2^n$	$\dots$
-2	4	$\dots$	$(-2)^n$	$\dots$

U that correctly represents their interval scale structure, so that o_1 is at least as good as o_2 if and only if $U(o_1) \geq U(o_2)$, as well as so that o_2 is a proportion of x of the way from o_1 to o_3 if and only if

$$U(o_2) = (1-x)U(o_1) + xU(o_3).$$

Expected Utility Theory says that, for any utility representation U, a prospect at least as great expected utility is at least as good. For any utility representation U, an affine permutation of outcomes performs some affine transformation of real numbers $ax+b$ on the utility of those outcomes, so the ordering of prospects by their expected utility is invariant under positive affine permutations of outcomes, and is reversed by negative affine permutations of outcomes.

However, Expected Utility Theory does not quite establish the affine symmetry principles because not every prospect has an expected utility. A good example, and one we will return to in Section 4, is the *alternating St Petersburg prospect*, which, fixing an arbitrary utility representation, yields an outcome with utility $(-2)^n$ with probability $1/2^n$, for every positive integer n (see Table 1). The expected utility of this prospect would be given by the sum

$$-2 \times \frac{1}{2} + 4 \times \frac{1}{4} - 8 \times \frac{1}{8} + \dots = -1 + 1 - 1 + \dots,$$

which alternates between -1 and 0 and never converges. The affine symmetry principles apply even to prospects that lack expected utilities; saying, for example, that replacing all the outcomes in alternating St Petersburg with outcomes that are twice as far from the outcome called "Cake" should preserve which is better out of that prospect and certainty of Cake. If, for example, Cake has utility 0 on the utility representation used to describe the alternating St Petersburg prospect, and we assume alternating St Petersburg is better than Cake, then so should the doubled alternating St Petersburg, which yields an outcome of utility $2 \times (-2)^n$ with probability $1/2^n$ (as opposed to utility $(-2)^n$; again holding the utility representation fixed). This is the sort of judgement that the affine symmetry principles secure.

The affine symmetry principles are the topic of this paper. We begin in Section 2 by contrasting them with closely related but importantly different assumptions in decision theory. Section 3 lays out an orthodox axiomatic framework for decision theory that forms the backdrop within which the affine symmetry principles will be studied. Section 4 investigates the consequences of the affine symmetry principles in this setting, including both general structural consequences, such as the vindication of Mark Colyvan's (2008) Relative Expectation Theory, as well as consequences about the comparisons of particular prospects. Most strikingly, they imply that the alternating St Petersburg prospect is exactly equally good as a sure utility $-1/2$, and that the "Pasadena" prospect of Nover and Hájek, 2004 is equally good as utility $\ln 2$, among other comparisons. Model theory for the affine symmetry principles is developed in Section 5, establishing among other things the consistency of the principles. The affine symmetry principles, then, form a well-motivated, powerful, and consistent way of extending decision theory in cases where expected utility is undefined.

We will assume throughout that outcomes are representable by real number utilities. This assumption rules out infinite ratios of differences between outcomes. There is no philosophical reason behind such a restriction, but it simplifies the mathematics significantly, so it is worthwhile to restrict our attention here to some domain of outcomes for which all ratios of differences are finite, so that the real number utility representation is possible.

2 Distinguishing the Affine Symmetry Principles

Lore in decision theory has it that utility is only 'defined' (or 'unique') up to a positive affine transformation. What this means is that for any utility representation U, any positive affine transformation of U (i.e., $aU + b$ for a positive real number a and real number b) represents the interval scale structure of outcomes equally well. This being the case, it can be hard to see why the affine symmetries, or at least Positive Affine Invariance, could add anything to standard decision theory. But there is an important difference: the lore is a well-established theorem of decision theory (originating in von Neumann and Morgenstern, 1944), whereas the affine symmetry principles are substantial additional axioms. Understanding the affine symmetry principles requires getting clear on the difference.

The difference between the lore and the affine symmetry principles can be illustrated by recognizing the difference between an affine transformation of a utility representation and an affine permutation of outcomes. A utility representation is a function from outcomes to numbers that represent those outcomes. An affine transformation of the utility representation is another function from outcomes to numbers, so that different numbers are chosen to represent the same outcomes—see Figure 3. We may think of the difference as a change in units, such as representing the freezing point of water by the number 0 (as in degrees Celsius) or the number 32 (as in degrees Fahrenheit). When it is observed that utility is 'defined' up to a positive affine transformation, what is being observed is that standardly, the role of a utility representation is to numerically represent the interval scale structure of outcomes, so that any positive affine transformation of the utility representation would do an equally good job in that role. By contrast, an affine permutation of outcomes is a function from outcomes to outcomes: outcomes themselves replace outcomes on an affine permutation of outcomes, and numbers need not enter the picture—see Figure 4. The theorem is that the interval scale structure of outcomes is faithfully represented by any positive affine transformation of a utility representation. The principle of Positive Affine Invariance adds that nothing besides the interval scale structure of outcomes enters into the comparison of prospects.[1]

The distinction is somewhat blurred in the literature, especially in connection with unbounded utility. In the following passage, Easwaran argues from the observation that utility is only defined up to a positive affine transformation to a principle much like Positive Affine Invariance:

> Hájek (2013, pp. 9–10) [Hájek, 2014: 541-42] notes that utility is only defined up to a shift and a stretch—there is no well-defined 0 and no well-defined unit in which utility is measured. Thus, it ought to be the case that adding a constant to the utility of every outcome of a gamble affects the overall value of the gamble by adding the same constant, and similarly for multiplying the utilities by a positive constant. (Easwaran, 2014b: 523)

The subtle point is that the principle being argued for is one that also refers to the utility representation, so it would seem that affine permutations of outcomes never enter the picture. But

[1] This distinction is closely related to the distinction between passive and active transformations in geometry—see page 84 of Struik, 1953.

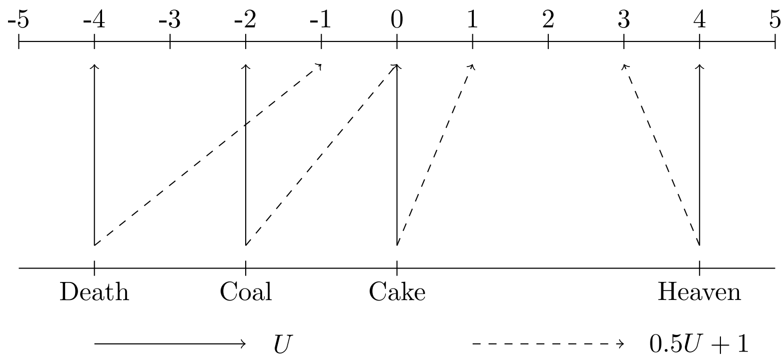

Figure 3: U and $0.5U + 1$ are utility representations, so are equally good representations of the interval scale structure of outcomes.

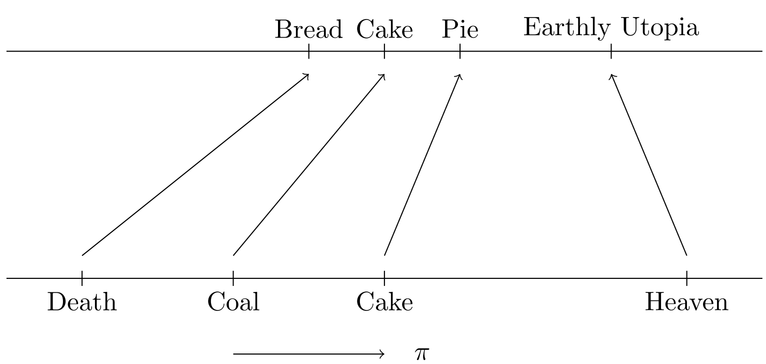

Figure 4: π is a positive affine transformation, because it preserves the ordering and ratios of differences between prospects (in relation to Figure 3, $\pi = U^{-1} \circ (0.5U + 1)$). Positive Affine Invariance says that substituting Death for Bread, Coal for Cake, etc. in two prospects preserves their comparison.

Easwaran is not arguing for a triviality. What follows 'it ought to be the case that' is not deductively entailed by what precedes 'thus'. Rather, the picture that Easwaran is working with is that a typical decision theory specifies an ordering of (some) prospects by an operation that converts a utility representation into a real-valued *value function* on prospects, so that when two prospects are given values in this sense, the one with the greater value is better. Symbolically, an Easwaran-style theory posits a two-place value function $V(\cdot,\cdot)$, which takes a utility representation and a prospect and outputs a number. Expected Utility Theory is a theory of this sort, where $V(U,\mu)$ is the expectation of U on μ, or $\mathbf{E}_\mu U$, as are the proposals investigated in Easwaran's paper. Easwaran's constraint amounts to forbidding such theories where the ordering of prospects achieved at the end is sensitive to which utility representation is plugged in at the beginning, which is to say that for any prospects μ and ν, any utility representation U, and positive real numbers a and real number b

$$V(U,\mu) \geq V(U,\nu) \leftrightarrow V(aU+b,\mu) \geq V(aU+b,\nu).$$ [2]

Easwaran's constraint is motivated by the same idea we have used to motivate the affine symmetry principles: a theory which violated the constraint would be sensitive to features of prospects other than probability and the interval scale structure of their possible outcomes; such a theory would posit an as-yet unrecognized feature of outcomes to be relevant in the evaluation of prospects, and it seems very odd to accept such a theory. But it is not something that is established just by noting that "utility is only defined up to a shift and a stretch."

Adopting Easwaran's constraint does not require accepting Positive Affine Invariance. Rather, the constraint forbids accepting theories which posit specific comparisons that contradict Positive Affine Invariance, such as a theory which gives value $-1/2$ to the alternating St Petersburg prospect but 3 to double the alternating St Petersburg prospect (this would be a counterexample to Positive Affine Invariance because double $-1/2$ is -1, so double the alternating St Petersburg would have to be no better than utility -1). Whereas Positive Affine Invariance says that theories violating Easwaran's constraint are false, Easwaran's constraint only forbids us from accepting them, which could be for any number of reasons (for example, it might be because we can't know any such theory to be true). Nevertheless, Positive Affine Invariance would be a very natural explanation of why adopting Easwaran's constraint is a good idea.

3 Axiomatic Decision Theory

The affine symmetry principles will be investigated within a variant of the classical system of von Neumann and Morgenstern (1944), with some standard modifications to accommodate the possibility of prospects which have infinitely many possible outcomes of arbitrarily great utility. We take as primitive a measurable space of *outcomes*, and define a *prospect* to be a probability distribution over outcomes, and use the greek letters μ, ξ, ν as variables ranging over prospects. We use '$\succsim$' for the *at least as good* relation between prospects which is the main topic of decision theory. $\succsim$ is assumed to be a *total preordering* of prospects; which is a reflexive and transitive relation that also compares every pair of prospects. We also assume von Neumann and Morgenstern's principle of *Independence*, which says that μ is at least as good as ν if and only if, for any third prospect ξ,

[2] Equivalently, given the Independence principle (Section 3) and if V extends expected utility:

$$V(aU+b,\mu) = aV(U,\mu) + b$$

an x chance of μ with a $1-x$ chance of ξ is at least as good as an x chance of ν with a $1-x$ chance of ξ. In symbols, $\mu \succsim \nu$ if and only if

$$x\mu + (1-x)\xi \succsim x\nu + (1-x)\xi.$$

Also in line with von Neumann and Morgenstern, we assume that outcomes have interval scale structure as outlined in the introduction. This amounts to the claim that there is a *utility representation* of outcomes. This is a measurable function U from outcomes to real numbers, such that for any prospects μ and ν that can only yield finitely many outcomes, the expectation of U on μ is at least as great as the expectation of U on ν. That is, when μ yields outcomes o_1 through o_n with probabilities x_1 through x_n, and ν yields outcomes u_1 through u_m with probabilities y_1 through y_m, we compare μ and ν by comparing

$$\mathbf{E}_\mu U = \int U \,\mathrm{d}\mu = x_1 U(o_1) + \cdots + x_n U(o_n), \text{ and}$$
$$\mathbf{E}_\nu U = \int U \,\mathrm{d}\nu = y_1 U(u_1) + \cdots + y_m U(u_m).$$

We assume also that any such utility representation is a bijection, which is to say that every real level of utility is exemplified by some outcome, and that equally good outcomes are identified.

The main divergence from von Neumann and Morgenstern is that we omit the Archimedean principle (which says that every prospect is equally good as some outcome),[3] in favour of the widely accepted principle of *Stochastic Dominance*.[4] A prospect μ is said to *strictly stochastically dominate* ν if for every outcome o, μ gives at least as great probability as ν does to getting an outcome better than o, and if for some o, μ gives a strictly greater probability of getting an outcome better than o. Stochastic dominance states that a prospect which strictly stochastically dominates another is strictly better. Intuitively, it says that modifying a prospect by shifting some probability of worse outcomes to better outcomes is a strict improvement. Combining Stochastic Dominance with the existence of a bijective utility representation implies that a St Petersburg prospect, which yields utility 2^n with probability $1/2^n$ for each $n \geq 1$, is better than any outcome, in contradiction with the Archimedean principle.

Putting these principles together yields the following theory, which following Goodsell (2024) we call DTU:

Total Preordering $\succsim$ is reflexive ($\mu \succsim \mu$), transitive (if $\mu \succsim \xi \succsim \nu$, then $\mu \succsim \nu$), and total ($\mu \succsim \nu$ or $\nu \succsim \mu$).

Independence $\mu \succsim \nu$ if and only if for any prospect ξ and number x between 0 and 1, $x\mu + (1-x)\xi \succsim x\nu + (1-x)\xi$.[5]

[3]More precisely, the Archimedean principle says that for any prospects $\mu \succsim \xi \succsim \nu$, some mixture of μ and ν is equally good as ξ: $x\mu + (1-x)\nu \sim \xi$ for some x. This is to be distinguished from an Archimedean principle for *outcomes*, which says the same when μ, ξ, ν are each 100% likely to yield some outcome. The Archimedean principle for outcomes follows from Simple Expected Utility Theory.

[4]In recent philosophical literature on unbounded utility, Stochastic Dominance is defended in Nover and Hájek, 2004, Easwaran, 2014a, Meacham, 2019, and Russell, forthcoming.

[5]Addition and multiplication of prospects are defined eventwise, so $x\mu + (1-x)\xi$ is the probability distribution that assigns to each measurable set X the probability

$$x\mu(X) + (1-x)\nu(X).$$

Simple Expected Utility Theory There is a bijective utility representation: a measurable bijection U from outcomes to real numbers such that for any simple μ and ν, $\mu \succsim \nu$ if and only if $\mathbf{E}_\mu U \geq \mathbf{E}_\nu U$.

Stochastic Dominance For any utility representation U, if

$$\mu(\{u : U(u) > U(o)\}) \geq \nu(\{u : U(u) > U(o)\})$$

holds for every outcome o, and the comparison is strict for some o, then $\mu \succ \nu$.

Various objections to DTU are possible.[6] However, DTU is incontestably a natural and orthodox approach to decision theory, and it is strong enough to bring out some of the interesting consequences of the affine symmetry principles. Importantly, the theory is on its own consistent, as Goodsell proves. It will be assumed in what follows.

A natural addition to DTU would be *expected utility theory*, which says that prospects can be compared by their expected utilities whenever they are defined, and not just when the prospects being compared are simple (as in Simple Expected Utility Theory). Goodsell (2024, Theorem 3) proves that DTU does not quite imply Expected Utility Theory. A very natural addition to DTU that fills this gap is the principle of L^1 Continuity, which says that the preordering $\succsim$ is continuous with respect to convergence in expectation of any utility representation. To be more precise, for a utility representation, U, the L^1 *U-distance* between prospects μ and ν is the infimal expectation of the absolute difference in U between any two random variables distributed according to μ and ν (the distance is ∞ if the expected difference diverges for all such random variables). L^1 Continuity may then be stated thus:

L^1 **Continuity** If for any utility representation U a sequence μ_i of prospects that converges in L^1 U-distance to μ, then if $\mu_i \succsim \nu$ for each i, $\mu \succsim \nu$ as well.[7]

Adding L^1 Continuity to DTU yields Expected Utility Theory, and also brings out some further interesting consequences of the affine symmetry principles, which will be noted in what follows.

To reason with affine permutations of outcomes, we first coordinatize the outcomes by fixing some arbitrary utility representation U, and adopt the notation of writing $aU + b$ for the permutation which maps an outcome of utility x to the outcome of utility $ax + b$. More accurately, this permutation would be denoted $U^{-1} \circ (aU + b)$ (see Figures 3 and 4); the benefit of the compact notation is that it allows us to write $[aU + b]$ for the function which modifies a prospect by replacing each outcome by the permuted outcome; which equivalently replaces any chance that prospect has of utility x with the same chance of utility $ax + b$. Permutations of this form with $a \neq 0$ are all and only the affine permutations of prospects (the choice of U doesn't matter). We may now state the affine symmetry principles as follows:

Positive Affine Invariance For any utility representation U, if $a > 0$, then $\mu \succsim \nu$ if and only if $[aU + b]\mu \succsim [aU + b]\nu$.

[6] Lauwers, 2016 shows that totality is not acceptable from a mathematical constructivist point of view, given the other assumptions. Savage (1954), Hammond (1998), and more recently Russell and Isaacs (2021) argue for a *bounded*, rather than bijective, utility representation, on the basis of principles related to Savage's P7. Moreover, the very idea of identifying prospects with probability distributions over outcomes is called into question by Seidenfeld et al., 2009 as well as Lauwers and Vallentyne, 2017, who argue that there can be differences in choiceworthiness between actions which would yield the same probability distributions over outcomes.

[7] Although L^1 U-distance depends on U, the topology generated does not.

Negative Affine Anti-Invariance For any utility representation U, if $a > 0$, $\mu \succsim \nu$ if and only if $[-aU + b]\mu \precsim [-aU + b]\nu$.

In fact, only Negative Affine Anti-Invariance is strictly needed, since it straightforwardly implies Positive Affine Invariance in DTU. However, for the purpose of understanding proofs, it is helpful to replace Negative Affine Anti-Invariance with the weaker Reflection Anti-Invariance:

Reflection Anti-Invariance For any utility representation U, $\mu \succsim \nu$ if and only if $[-U]\mu \succsim [-U]\nu$.

Positive Affine Invariance is independent of Reflection Anti-Invariance, but together these principles entail and are entailed by Negative Affine-Invariance (Reflection Invariance does not have an obvious interpretation in terms of the interval scale structure of outcomes, but is easy to use once a utility representation has been decided upon). DTU+Sym will be DTU plus the affine symmetry principles.

4 Consequences of Affine Symmetry

We now turn to the question of how adding the affine symmetry principles affects decision theory in DTU. The most important results are that the affine symmetry principles cannot be derived or refuted in DTU:

Theorem 1. *Neither affine symmetry principle is a theorem of* DTU*. The same is true when* DTU *extended by any or all of L^1 Continuity, Relative Expectation Theory (page 10; see also Colyvan, 2008), Weak Expectation Theory (Easwaran, 2008), or Principal Value Theory (Easwaran, 2014b).*

Proof. Goodsell, 2024 (see Remark 3 for a countermodel to Reflection Anti-Invariance and Theorem 10 for a countermodel to Positive Affine Invariance). □

Theorem 2. *The affine symmetry principles are consistent with* $\mathsf{DTU} + L^1$ *Continuity.*

Proof. Section 5. □

In light of these two theorems, the affine symmetry principles are a nontrivial but very natural way of extending to DTU. In what follows we will build a picture of what they add to DTU by establishing some consequences of DTU + Sym that are independent of DTU (proofs of independence are omitted, but can be easily generated by modifying examples from Goodsell, 2024). The overarching theme will be that adding the affine symmetry principles to DTU vindicates and unifies much previous literature on decision theory with unbounded utility, especially the work of Colyvan (2008) and to a lesser extent Easwaran (2008; 2014b) and Hájek (2014).

To begin, the affine symmetry principles do not entail that a doubling of utility yields a doubling of overall value, in the sense that $\mu \sim 0.5\delta_0 + 0.5[2U]\mu$ where δ_x is the prospect that yields utility x with 100% probability; this identity of value provably fails when μ is a St Petersburg prospect, because then we have $\mu = 0.5\delta_2 + 0.5[2U]\mu$.

By contrast, they do imply that adding to a fixed amount of utility to each outcome adds to the overall value of a prospect, in the sense that fifty-fifty between some prospect μ and and utility b, and fifty-fifty between that prospect μ plus b units of utility ($[U + b]\mu$) and utility 0, are equally good. More generally:

Theorem 3 ($\mathsf{DTU} + \mathrm{Sym}$)**.**

$$x\left[U+\frac{b}{x}\right]\mu+(1-x)\nu \sim x\mu+(1-x)\left[U+\frac{b}{1-x}\right]\nu. \tag{1}$$

Proof. By Reflection Anti-Invariance we have for any prospect ξ and constant c,

$$\frac{1}{2}\left[U+\frac{c}{2}\right]\xi+\frac{1}{2}\left[-U-\frac{c}{2}\right]\xi \sim \delta_0.$$

Shifting utility up by $c/2$ yields

$$\frac{1}{2}\left[U+c\right]\xi+\frac{1}{2}[-U]\xi \sim \delta_{c/2}. \tag{2}$$

Moreover, (1) is equivalent to

$$\delta_0 \sim \frac{x}{2}\left[U+\frac{b}{x}\right]\mu+\frac{1-x}{2}\nu+\frac{x}{2}[-U]\mu+\frac{1-x}{2}\left[-U-\frac{b}{1-x}\right]\nu$$

The right hand side of which can be divided into two parts which have the form of (2), whence the equality of value may be derived with Simple Expected Utility Theory:

$$\delta_0 \sim x\left(\frac{1}{2}\left[U+\frac{b}{x}\right]\mu+\frac{1}{2}[-U]\mu\right)+(1-x)\left(\frac{1}{2}\nu+\frac{1}{2}\left[-U-\frac{b}{1-x}\right]\nu\right).$$

□

This result can be extended, in the presence of L^1 Continuity, to vindicate the central principle of Colyvan, 2008, *Relative Expectation Theory*. Stating Relative Expectation Theory in the present setting, where prospects are identified with probability distributions, is a little awkward. Here, Colyvan's principle is stated by relating the prospects to be compared with random variables on $[0,1]$ that are distributed according to those prospects.[8]

Relative Expectation Theory Let X and Y be any outcome-valued random variables on $[0,1]$ with Lebesgue measure that are distributed according to μ and ν respectively. Let $X-Y$ be the outcome-valued random variable on $[0,1]$ that yields utility $x-y$ when X yields utility x and Y yields utility y. Then if $X-Y$ has an expected utility, $X \succsim Y$ if and only if the expected utility of $X-Y$ is greater than or equal to 0. See Figure 5.

Like how Expected Utility Theory is distinct from Simple Expected Utility Theory when L^1 Continuity is not assumed, we also distinguish a version of Relative Expectation Theory that applies only when the difference of the relevant random variables is simple:

Simple Relative Expectation Theory Let X and Y be outcome-valued random variables on $[0,1]$ with Lebesgue measure that are distributed according to μ and ν respectively. Let $X-Y$ be the outcome-valued random variable on $[0,1]$ that gives yields outcome $x-y$ when X yields x and Y yields y. Then if $X-Y$ is simple (i.e., has finite codomain), then $X \succsim Y$ if and only if the expected utility of $X-Y$ is greater than or equal to 0.

[8]The principle is more natural in a setting where prospects are identified with random variables over an atomless

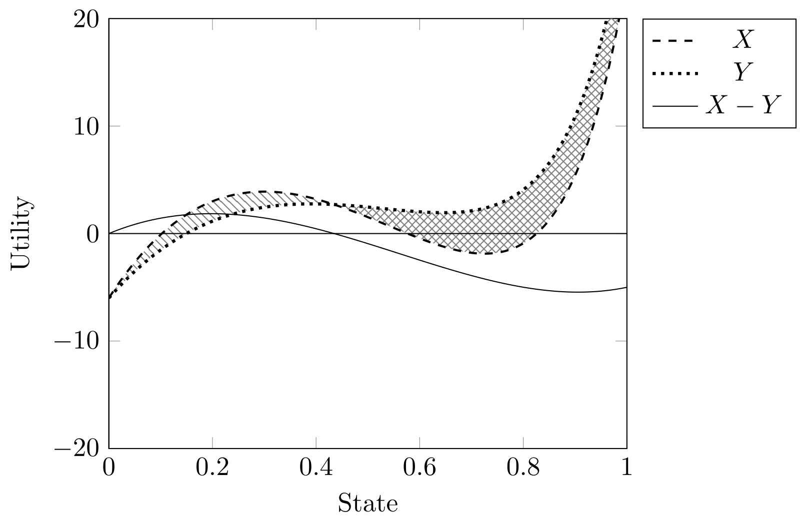

Figure 5: Graphs of the utility of outcome-valued random variables on $[0, 1]$, X and Y, and their difference $X - Y$. Relative Expectation Theory says that if $X - Y$ has an expected utility, then the distribution of X is at least as good as the distribution of Y just in case $X - Y$ has a nonnegative expected utility. Equivalently: if the hatched and crosshatched areas both have finite area, then the distribution of X is better than the distribution of Y just in case the northwest-hatched shape has an area at least as great as the crosshatched shape. Simple Relative Expectation Theory says the same in the case where $X - Y$ is simple (which is not so in the figure).

Both principles are quite the mouthful, and may not seem particularly natural in this setting. This does not matter, since both can be derived from very natural assumptions:

Theorem 4 ($\mathsf{DTU} + \mathrm{Sym}$). *Simple Relative Expectation Theory is true.*

Proof. First note that $\mu \succsim \nu$ is equivalent to $0.5\mu + 0.5\delta_0 \succsim 0.5\nu + 0.5\delta_0$. Let X_μ, X_ν, and X_ξ be distributed according to $0.5\mu + \delta_0$, $0.5\nu + 0.5\delta_0$, and $0.5\xi + 0.5\delta_0$ respectively such that $U \circ X_\mu = U \circ X_\nu + U \circ X_\xi$. Let o be one of the outcomes of ξ, and let X_μ^1 modify X_μ by subtracting o from all those outcomes where X_ξ yields o, and adding the same probability of o to the region where X_μ yields utility 0. By Theorem 3, the distribution of X_μ^1 is equally good as μ. Repeating this process for the n possible outcomes of ξ, we find the distribution of X_μ^n is equally good as μ, and the distribution of X_μ^n is $0.5\nu + 0.5\xi$, whence the result follows by Independence and Simple Expected Utility Theory. □

Corollary 5 ($\mathsf{DTU} + \mathrm{Sym}$). *Relative Expectation Theory is equivalent to L^1 Continuity.*

We will say that μ differs from ν by ξ when for some X and Y distributed according to μ and ν, $X - Y$ is distributed according to ξ.

Notice that by Reflection Anti-Invariance, a symmetric prospect—a prospect μ that is equal to $[-U]\mu$—is equally good as zero utility. From Relative Expectation Theory it also follows that if μ and $[-U]\mu$ differ by a prospect of finite expectation ξ, then μ differs by the expectation of ξ from a symmetric prospect, so μ is equally good as the expected utility of ξ. From this observation, we may derive a useful strengthening of Expected Utility Theory that does not seem to have appeared in the literature. For a prospect μ, let F be the *cumulative distribution function* of μ, which maps each real number x to the probability that μ yields an outcome with utility at least x. When μ has an expected utility, it is given by the integral

$$\mathbf{E}_\mu U = \int_0^\infty F(z)\,\mathrm{d}z - \int_{-\infty}^0 1 - F(z)\,\mathrm{d}z$$

Set $F^+(z) = F(z)$, and set $F^-(z) := 1 - F(-z)$. Then when the expected utility of μ exists, we also have

$$\int_0^\infty F(z)\,\mathrm{d}z - \int_{-\infty}^0 1 - F(z)\,\mathrm{d}z = \int_0^\infty F^+(z) - F^-(z)\,\mathrm{d}z$$

but the right-hand-side integral might exist even when the expectation does not. For example, if μ is symmetric but lacks an expectation, then $F^+ = F^-$, so the right-hand-side integral exists and is zero. Let the integral of $F^+ - F^-$ from 0 to infinity be the *folded expected utility* of μ.

Folded Expectation Theory If μ and ν have folded expected utilities (see Figure 6), then $\mu \succsim \nu$ if and only if the folded expected utility of μ is at least as great as that of ν.

Theorem 6 ($\mathsf{DTU} + \mathrm{Sym}$). *Folded Expectation Theory follows from L^1 Continuity.*

Proof. Let μ be a prospect with cumulative distribution function F. Consider the random variables X^+, defined to agree with $(F^+(z))^{-1}$ almost everywhere where $F(z)$ is positive, and to be identically zero otherwise, X^-, which is the same but for F^-, and their difference $X^+ - X^-$ (see Figure 7).

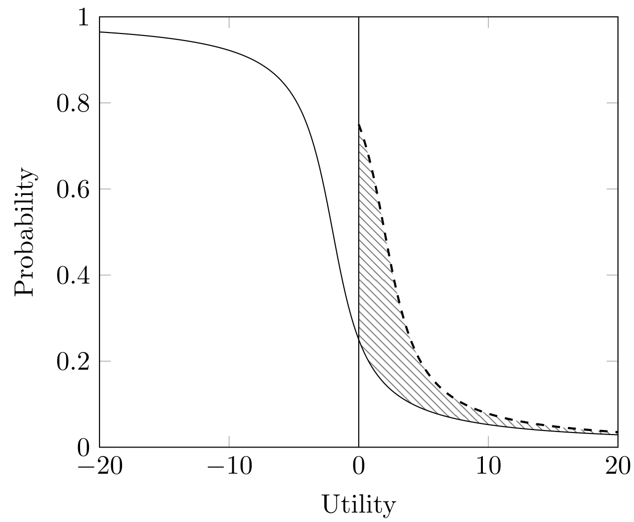

Figure 6: If the cumulative distribution function over utility for a prospect is the solid curve, then the folded expected utility of that prospect is equal to negative one times the hatched area, if that area is finite.

If μ has a folded expected utility, then the distribution of $X^+ - X^-$ has an expected utility equal to the folded expected utility of μ, so μ is equally good as its folded expected utility by Theorem 5. □

Folded Expectation Theory should be distinguished from some related proposals in the literature. It is (on its own) incomparable in strength with both Easwaran's Weak Expectation Theory (Easwaran, 2008) and the stronger Principal Value Theory (Easwaran, 2014b). It is not weaker since unlike these two proposals, Folded Expectation Theory evaluates every symmetric prospect as equally good as utility 0; Easwaran's two principles sometimes fail to assign values to such prospects. It is not stronger than either principle for a more complicated reason. A prospect μ has a weak expectation when the following two conditions are met. First, the definite integral

$$\int_0^t F^+(z) - F^-(z)\, \mathrm{d}z$$

must have a limit as t goes to infinity; which is the *principal value* of the integral. Second, μ must have 'thin tails': the probability that μ yields an outcome better than utility t or worse than $-t$ must decay with $o(t^{-1})$ as t goes to infinity. In this case the *weak expected utility* of the prospect is given by the principal value of the integral. These conditions are met when μ is the *Pasadena* prospect, which for each positive integer n yields utility $-(-2)^n$ with probability $1/2^n$ (see Table 2). Easwaran shows the Pasadena prospect to have a weak expected utility of $\ln 2$. By contrast, the prospect does not have a folded expected utility.

Proposition 7. *The Pasadena prospect does not have a folded expected utility.*

probability space of "states", in the style of Savage, 1954.

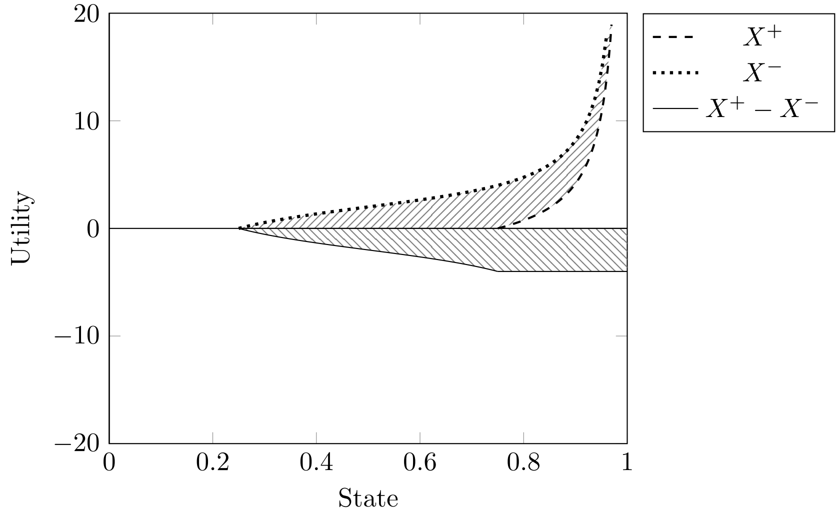

Figure 7: Graphs of the random variables X^+, X^-, and $X^+ - X^-$ from Theorem 6, in the case where F is the cumulative distribution function from Figure 6. Notice that the two hatched shapes have equal area, and equal area to the hatched shape in Figure 6.

Table 2: The Pasadena prospect.

$1/2$	$1/4$	$\dots$	$1/2^n$	$\dots$
2	$-4/2$	$\dots$	$-(-2)^n/n$	$\dots$

Table 3: The Arroyo prospect.

$1/2$	$1/6$	$\ldots$	$1/(n^2+n)$	$\ldots$
2	-3	$\ldots$	$(-1)^{n+1}(n+1)$	$\ldots$

Table 4: The Arroyo prospect + 0.5 units of utility.

$1/2$	$1/6$	$\ldots$	$1/(n^2+n)$	$\ldots$
2.5	-2.5	$\ldots$	$(-1)^{n+1}(n+1)+0.5$	$\ldots$

For a nontrivial application of Folded Expectation Theory, we may evaluate Bartha's (2016) *Arroyo* prospect (Table 3), which for each positive integer n yields utility $(-1)^{n+1}(n+1)$ with probability $1/(n^2+n)$. Like the Pasadena prospect, the Arroyo prospect's expected utility would be given by the divergent sum

$$1 - 1/2 + 1/3 + \ldots,$$

which conditionally converges to $\ln 2$. Bartha demonstrates that the Arroyo prospect does not have a weak expected utility. But it does have a folded expected utility of $\ln 2$:

Theorem 8 (DTU + Sym + L^1 Continuity)**.** *The Arroyo prospect is equally good as utility* $\ln 2$.

Proof. Let A be the Arroyo prospect, and let F be the cumulative distribution function for $[U{+}0.5]A$ (Table 4). Then $F^+(z) - F^-(z)$ is the cumulative distribution function of a prospect which yields utility $2n + 0.5$ with probability

$$\frac{1}{2n(2n-1)} - \frac{1}{2n(2n+1)},$$

and such a prospect has expected utility $0.5 + \ln 2$. □

DTU + Sym, with or without L^1 Continuity, also goes beyond Folded Expectation Theory in many ways. We shall find that the Pasadena prospect, unevaluable by Folded Expectation Theory, must receive utility $\ln 2$ on DTU + Sym + L^1 Continuity. To achieve this result, we will first give a simpler example of some the necessary techniques by evaluating the alternating St Petersburg prospect from Section 1 (Table 1), which yields utility $(-2)^n$ with probability $1/2^n$ for each positive integer n. The alternating St Petersburg paradox does not have folded expected utility, or a weak expectation, or a principal value (the relevant integral oscillates between -1 and 0). Nevertheless, we find:

Theorem 9 (DTU + Sym)**.** *The alternating St Petersburg prospect has utility* $-1/2$.

Proof. Let P be the alternating St Petersburg prospect. We have the identity

$$P = 0.5\delta_{-2} + 0.5[-2U]P.$$

Suppose $P \succ \delta_{-1/2}$. By the affine symmetry principles, it follows that $[-2U]P \prec \delta_1$. By Simple Expected Utility Theory and Independence, we have

$$0.5\delta_{-2} + 0.5[-2U]P \prec 0.5\delta_{-2} + 0.5\delta_1 \sim \delta_{-1/2}$$

so $P \prec \delta_{-1/2}$, a contradiction. An analogous contradiction arises from supposing $P \succ \delta_{-1/2}$, so by Totality we have $P \sim \delta_{-1/2}$. □

This result can be generalized to any prospect that is self-similar in the sense of being related to itself by a convex combination with an arbitrary prospect ν and itself under a negative affine transformation:

Theorem 10 (DTU+Sym)**.** *Suppose $\mu \sim (1-x)\nu + x[-aU+b]\mu$, with $a > 0$ and $1 > x > 0$. Then*

$$\mu \sim \sum_{n=0}^{\infty}(1-x)x^n \left[-a^n U + \frac{(-a)^n bx}{1-x}\right] \nu.$$

Proof. Define

$$E\mu := (\mu \sim (1-x)\nu + x[-aU+b]\mu),$$

notice first that by Theorem 3, $E\mu$ is equivalent to

$$\mu \sim (1-x)\left[U + \frac{bx}{1-x}\right]\nu + x[-aU]\mu.$$

We now show that solutions to E are equally good; which is to say that if $E\mu$ and $E(\mu')$ then $\mu \sim \mu'$. Suppose for contradiction that $E\mu$ and $E\mu'$ and $\mu \succ \mu'$. Then, by the affine symmetries and independence,

$$(1-x)\left[U + \frac{bx}{1-x}\right]\nu + x[-aU]\mu \prec (1-x)\left[U + \frac{bx}{1-x}\right]\nu + x[-aU]\mu',$$

so $\mu \prec \mu'$, a contradiction.

It remains only to show that $E\mu_E$, where

$$\mu_E := \sum_{n=0}^{\infty}(1-x)x^n \left[-a^n U + \frac{(-a)^n bx}{1-x}\right] \nu.$$

for this, it suffices to expand the definition as follows:

$$\begin{aligned}\mu_E &= (1-x)\left[U + \frac{bx}{1-x}\right]\nu + \sum_{n=1}^{\infty}(1-x)x^n \left[-a^n U + \frac{(-a)^n bx}{1-x}\right] \nu \\ &= (1-x)\left[U + \frac{bx}{1-x}\right]\nu + x[-aU]\mu_E,\end{aligned}$$

so $E\mu_E$ follows by reflexivity of $\sim$. □

Theorem 10 yields that a self-similarity of value of the form

$$\mu \sim (1-x)\nu + x[aU+b]\mu$$

has a unique solution in value when $a < 0$; which must therefore be the value of the unique prospect which solves the self-similarity equation of prospects

$$\mu = (1-x)\nu + x[aU+b]\mu$$

Table 5: The Highland Park prospect, H.

$1/2$	$1/4$	$1/8$	$1/16$	$\dots$	$1/2^{2n-1}$	$1/2^{2n}$	$\dots$
2	-4	$8/3$	$-16/3$	$\dots$	$2^{2n-1}/(2n-1)$	$-2^{2n}/(2n-1)$	$\dots$

It is evident that the uniqueness of solution to the equation of value no longer holds when $a > 0$. For let $a = 1/x$, and $\nu = \delta_{(-xb)/(1-x)}$, and the resulting equation of value,

$$\mu \sim (1-x)\delta_{\frac{-xb}{1-x}} + x\left[\frac{1}{x}U + b\right]\mu$$

is solved for any simple μ.

Let us now return to the Pasadena prospect (Table 2). The series that would give its expected utility,

$$\sum_{n=1}^{\infty} \frac{1}{2^n} \frac{-(-2)^n}{n}$$

diverges, but the sequence of its partial sums in the given order,

$$1 - \frac{1}{2} + \frac{1}{3} - \dots$$

conditionally converges to $\ln 2$. It is natural to suppose that the Pasadena prospect really is equally good as utility $\ln 2$, and this judgement has been supported by a variety of principles in the literature, including Weak Expectation Theory. Hájek (2014: 550), identifies five additional proposals as entailing this judgement, an observation he has called a 'hexangulation' argument for the Pasadena prospect being equally good as utility $\ln 2$. As previously mentioned, Hájek independently comes close to endorsing the affine symmetry principles (at the very least he seems willing to apply Easwaran's constraint—see Section 2—to rule out any theory which gives rise to specific counterexamples to the affine symmetry principles). It turns out that these two natural and seemingly unrelated opinions are not logically independent: the affine symmetry principles already imply a value of $\ln 2$ for the Pasadena prospect, given L^1 Continuity:

Theorem 11 (DTU+Sym+L^1 Continuity)**.** *The Pasadena prospect is equally good as the outcome of utility* $\ln 2$*.*

Proof. Consider a modification of the Pasadena prospect, the *Highland Park prospect* of Table 5, H, which sours the negative-utility-outcomes of the prospect by decreasing the denominator by 1. The Highland Park prospect worsens the Pasadena prospect by a prospect with a finite expectation; specifically the prospect

$$\frac{2}{3}\delta_0 + \frac{1}{3}\sum_{n=1}^{\infty} \frac{3}{2^{2n}} \delta_{\frac{-2^{2n}}{2n(2n-1)}}$$

which has expectation $-\ln 2$. By Relative Expectation Theory, the Highland Park prospect is $\ln 2$ units of utility worse than the Pasadena prospect.

Now, notice that the Highland Park prospect H has the following property

$$H = \frac{2}{3}Q + \frac{1}{3}[-2U]Q,$$

where

$$Q := \sum_{n=1}^{\infty} \frac{3}{2^{2n}} \delta_{\frac{2^{2n-1}}{2n-1}}.$$

Notice that Q and

$$\frac{3}{4}\delta_0 + \frac{1}{4}[4U]Q$$

also differ by a prospect of finite expectation, namely

$$\frac{3}{4}\delta_2 + \frac{1}{4}\sum_{n=1}^{\infty} \frac{3}{2^{2n-1}} \delta_{\frac{-2^{2n+1}}{(2n+1)(2n-1)}}$$

which has an expectation of 0, so

$$Q \sim \frac{3}{4}\delta_0 + \frac{1}{4}[4U]Q$$

and thus

$$Q \sim \frac{1}{2}\delta_0 + \frac{1}{2}[2U]Q$$

by Totality and Positive Affine Invariance. Then, by Reflection Invariance, H is equally good as 0. Therefore, since H worsens the Pasadena prospect by $\ln 2$ units of utility, the Pasadena prospect is equally good as the outcome of utility $\ln 2$. □

A natural conjecture is that $\mathsf{DTU}+\mathrm{Sym}+L^1$ Continuity entails one of Easwaran's Weak Expectation Theory or Principal Value Theory in addition to mere Folded Expectation Theory. Theorem 11 lends some plausibility to this conjecture, which is left for further work.

5 Models of Affine Symmetry

The main purpose of this section is to establish the consistency of $\mathsf{DTU} + \mathrm{Sym} + L^1$ Continuity (Theorem 2). This will be done by constructing a model of the theory. Future investigation of models of this sort will lead to better understanding of the limits of the theory.

A harmless simplification employed in the models will be the identification of the outcomes with the real numbers $\mathbf{R}$, with the real number x representing the outcome of utility x. Prospects may then be identified with the real-valued probability distributions. However, it will be useful to extend the convex space of probability distributions on $\mathbf{R}$ to the real vector space of *(finite) signed measures* (see, e.g., Folland, 1999 Chapter 3).

Definition 1 (Finite signed measure/$\mathcal{F}$)**.** A *finite signed measure* on $\mathbf{R}$ is a function on the measurable sets of $\mathbf{R}$ that is the difference $\mu^+ - \mu^-$ of finite measures μ^+ and μ^-. $\mathcal{F}$ is the vector space of finite signed measures.

The probability distributions are the positive signed measures which assign 1 to $\mathbf{R}$; the other signed measures have no decision-theoretically significant interpretation, and may be thought of as imaginary sums and differences of prospects where probabilities are allowed to be greater than 1 or less than 0. In this setting, principles of decision theory like Independence and Stochastic Dominance still make sense when applied to imaginary prospects, and assuming this additional structure makes the models easier to work with. In particular, we may use the standard identification of linear orderings with *cones* in the vector space (see, e.g., Aliprantis and Tourky, 2007 Chapter 1).

Definition 2 (Cone/convex cone/strict cone)**.** A (weak) *cone* is a set $C \subseteq V$ such that $x\mu \in C$ when $\mu \in C$ and $x > 0$.

A cone is *convex* when $x\mu + y\nu \in C$ whenever $\mu, \nu \in C$ and $x, y > 0$.

A *strict* cone is the set $C \setminus (-C)$, where C is a (weak) cone. When C is a cone, we write C^+ for the corresponding strict cone $C \setminus (-C)$, C^0 for the difference $C \setminus C^+ (= C \cap (-C))$, and C^- for $(-C^+)$.

A cone corresponds to a preordering $\geq_C$ of V, where $\mu \geq_C \nu$ iff $\nu \in \mu + C$, and $\mu >_C \nu$ iff $\nu \in \mu + C^+$, and $\mu =_C \nu$ iff $\nu \in \mu + C^0$.

When the cone is convex, the preordering $\geq_C$ is *linear*, in the sense that it satisfies for any positive reals x, y,

$$\mu \geq_C \nu \leftrightarrow x\xi + y\mu \geq_C x\xi + y\nu,$$

which is analogous to the Independence principle generalized to settings where probabilities needn't lie between 0 and 1.

We will build models of the affine symmetry principles by interpreting '$\succsim$' as the preordering of prospects determined by a cone $C_{\succsim}$ on $\mathcal{F}$. The principles governing '$\succsim$' are satisfied whenever the cone $C_{\succsim}$ has the following corresponding properties:

Totality $\mathcal{F} = C \cup C^-$.

Preordering $\mathbf{0} \in C$ (reflexivity), C is convex (transitivity).

Independence C is convex.

Expected Utility Theory and L^1 Continuity If $\int x \, \mathrm{d}\mu(x) \geq 0$ converges, then $\mu \in C$ if and only if $\int x \, \mathrm{d}\mu(x) \geq 0$.

Stochastic Dominance If $\mu([x, \infty)) \geq \nu([x, \infty))$ for all $x \in \mathbf{R}$, then $\mu - \nu \in C$. If in addition, $\mu([x, \infty)) > \nu([x, \infty))$, then $\mu - \nu \in C^+$.

The affine symmetry principles are decomposed into the following three constraints:

Scale Invariance When $a > 0$, $[aU]C \subseteq C$ and $[aU]C^+ \subseteq C^+$.

Reflection Anti-Invariance $[-U]C = -C$.

Prospect Shift Invariance When μ is a prospect (a probability distribution), $[U+b]\mu - (\mu + \delta_b) \in C^0$

In the setting of preorderings on $\mathcal{F}$ rather than the space of prospects, there is an important difference between scaling and shifting utility. Consider the imaginary prospect $\mathbf{0}$ which gives zero probability to every set. It is to be equally good as δ_0, which gives 100% probability to utility 0.

But improving the payoffs of $\mathbf{0}$ by a unit returns $\mathbf{0}$ again, which is worse than δ_0 so-shifted, which is δ_1. So, whereas we will be considering Scale Invariance as applying to all signed measures in $\mathcal{F}$, the analogous principle to Shift Invariance is restricted to prospects.

It is straightforward to come up with cones that satisfy these constraints besides Totality. One example is the cone that determines the continuous eventual dominance of truncated expectation ordering $\sqsupseteq$ of Lauwers (2016) and Goodsell (2024), defined thus:

$$\mu \sqsupseteq \nu := (\forall \varepsilon > 0)(\exists t > 0)(\forall s > t)\left(\int_{-s}^{s} x \, \mathrm{d}\mu - \int_{-s}^{s} x \, \mathrm{d}\nu > -\varepsilon\right)$$

$C_{\sqsupseteq} = \{\mu - \nu : \mu \sqsupseteq \nu\}$ is the corresponding cone. $C_{\sqsupseteq}$ can also be expressed as the set of signed measures whose principal value is eventually always greater than any negative number:

$$\mu \in C_{\sqsupseteq} = (\forall \varepsilon > 0)(\exists t > 0)(\forall s > t)\left(\int_{-s}^{s} x \, \mathrm{d}\mu(x) > -\varepsilon\right)$$

Theorem 12. *The cone $C_{\sqsupseteq} = \{\mu - \nu : \mu \sqsupseteq \nu\}$ satisfies the constraints besides Totality.*

Proof. Goodsell (2024). □

Note that $C_{\sqsupseteq}$ is not the only such cone, but it is a natural one with plausible decision-theoretic significance, and is easy to work with.

A model of the theory, then, can be generated by *extending* such a cone satisfying the constraints besides Totality, such as $C_{\sqsupseteq}$, to a cone that does have Totality while respecting the other constraints.

Definition 3 (Extension of a cone)**.** A cone D is said to (weakly) *extend* C if $C \subseteq D$ and $C^+ \subseteq D^+$; if $D \neq C$ then D is said to *strictly* extend C.

Some of the constraints are automatically had by any extension of a cone that satisfies them. These are the principles which specify exactly which elements a cone must have to satisfy them: Simple Expected Utility Theory, Stochastic Dominance, Prospect Shift Invariance, the reflexivity part of Preordering ($\mathbf{0} \in C$), and Totality (because there are no strict extensions to a cone which is already total). The other principles, Independence, Scale Invariance, and Reflection Anti-Invariance, constrain how the cone must relate to itself, and so will not be had by all extensions of a cone that has them.

Our strategy for constructing a cone satisfying the constraints including Totality, then, will be as follows: beginning with a cone like $C_{\sqsupseteq}$ that satisfies the constraints besides Totality, we will show that it is always possible to strictly extend this cone while still respecting the constraints of Independence, Scale Invariance, and Reflection Anti-Invariance. By Zorn's lemma, it will follow that there is an inextensible cone satisfying these principles, which will satisfy Totality.

Theorem 13 (Extensibility)**.** *Let C be a non-total convex cone satisfying Scale Invariance and Reflection Anti-Invariance. Then there is a strict extension of C satisfying those principles as well.*

Proof. Take $\mu \notin C \cup C^-$. Let C_μ be the smallest scale invariant and reflection anti-invariant convex cone containing μ:

$$C_\mu = \left\{\sum_{i=1}^{n} \operatorname{sgn}(a_i) x_i [a_i U]\mu : n \geq 1, x_i > 0, a_i \neq 0\right\}$$

We first show that C_μ intersects at most one of C^+, C^0, or C^-. Suppose without loss of generality the transformation

$$T_x^a := \sum_{i=1}^{n} \operatorname{sgn}(a_i) x_i [a_i U]$$

maps μ to an element of C^+ (otherwise substitute C^+ for C^0 or C^-), and consider an arbitrary transformation

$$T_y^b := \sum_{i=1}^{m} \operatorname{sgn}(b_i) y_i [b_i U].$$

$T_y^b(T_x^a \mu) \in C^+$ by the hypothesis that C is convex, scale invariant, and reflection anti-invariant. Notice that the two operations commute: $T_x^a \circ T_y^b = T_y^b \circ T_x^a$, so $T_x^a(T_y^b \mu) \in C^+$ as well. Therefore $T_y^b \mu$ cannot be in either C^0 or C^- since those two sets are scale invariant, reflection anti-invariant, convex, and disjoint from C^+.

Let D be $C + (C_\mu \cup \{\mathbf{0}\})$ if C_μ intersects C^+, $C + C_\mu + (-C_\mu)$ if C_μ intersects C^0, or $C + (-C_\mu \cup \{\mathbf{0}\})$ if C_μ intersects C^-, otherwise let D be $C + (C_\mu \cup \{\mathbf{0}\})$ (in effect supposing μ to be good when we could consistently suppose it to be any of good, bad, or neutral). Since D is the sum of scale invariant and reflection anti-invariant convex cones, D is also a scale invariant and reflection anti-invariant convex cone.

It is clear that $C \subseteq D$ and $C \neq D$. Now, suppose that $D = C + (C_\mu \cup \{\mathbf{0}\})$, then C_μ does not intersect $-C$, so $-C - (C_\mu \cup \{\mathbf{0}\})$ does not intersect C^+, so $C^+ \subseteq D^+$. Similarly in the other two cases. So $C^+ \subseteq D^+$, so D is a strict extension of C, as required. □

Remark 1. *The feature of rescaling and negating utility that ensures their extensibility is that they commute with the vector space operations (i.e., they are linear) and also commute with each other. This technique can therefore be used to establish the consistency of* DTU *with the monotonicity (with respect to $\succsim$) of any such class of operations, so long as that class is monotone with respect to $\sqsupseteq$.*

This establishes the consistency of $\mathsf{DTU} + \text{Sym}$. For the consistency of L^1 Continuity, it suffices to note that any extension of $C_\sqsupseteq$ has this property, since when μ and ν have a finite distance in the L^1 metric, $\mathbf{E}_{\mu-\nu}U$ is finite so $\mu - \nu$ is ranked as having that utility on $\sqsupseteq$.

References

Aliprantis, C. D., & Tourky, R. (2007). *Cones and duality*. American Mathematical Society. https://doi.org/https://doi.org/10.1090/gsm/084

Bartha, P. (2016). Making do without expectations. *Mind*, *125*(499), 799–827. https://doi.org/10.1093/mind/fzv152

Colyvan, M. (2008). Relative expectation theory. *Journal of Philosophy*, *105*(1), 37–44. https://doi.org/10.5840/jphil200810519

Easwaran, K. (2008). Strong and weak expectations. *Mind*, *117*(467), 633–641. https://doi.org/10.1093/mind/fzn053

Easwaran, K. (2014a). Decision theory without representation theorems. *Philosophers' Imprint*, *14*.

Easwaran, K. (2014b). Principal values and weak expectations. *Mind*, *123*(490), 517–531. https://doi.org/10.1093/mind/fzu074

Folland, G. B. (1999). *Real analysis: Modern techniques and their applications* (Vol. 40). John Wiley & Sons.

Goodsell, Z. (2024). Decision theory unbound. *Noûs*, *58*(3), 669–695. https://doi.org/10.1111/nous.12473

Hájek, A. (2014). Unexpected expectations. *Mind*, *123*(490), 533–567. https://doi.org/10.1093/mind/fzu076

Hammond, P. (1998). Objective expected utility. In S. Barbera, P. Hammond, & C. Seidl (Eds.), *Handbook of utility theory* (pp. 145–212, Vol. 1). Dordrecht: Kluwer Academic Publishers.

Lauwers, L. (2016). Why decision theory remains constructively incomplete. *Mind*, *125*(500), 1033–1043. https://doi.org/10.1093/mind/fzv174

Lauwers, L., & Vallentyne, P. (2017). A tree can make a difference. *Journal of Philosophy*, *114*(1), 33–42. https://doi.org/10.5840/jphil201711412

Meacham, C. J. G. (2019). Difference minimizing theory. *Ergo: An Open Access Journal of Philosophy*, *6*. https://doi.org/10.3998/ergo.12405314.0006.035

Nover, H., & Hájek, A. (2004). Vexing expectations. *Mind*, *113*(450), 237–249. https://doi.org/10.1093/mind/113.450.237

Russell, J. S. (forthcoming). Fixing stochastic dominance. *The British Journal for the Philosophy of Science*. https://doi.org/10.1086/728716

Russell, J. S., & Isaacs, Y. (2021). Infinite prospects. *Philosophy and Phenomenological Research*, *103*(1), 178–198. https://doi.org/10.1111/phpr.12704

Savage, L. (1954). *The foundations of statistics*. Wiley Publications in Statistics.

Seidenfeld, T., Schervish, M., & Kadane, J. (2009). Preference for equivalent random variables: A price for unbounded utilities. *Journal of Mathematical Economics*, *45*, 329–340. https://doi.org/https://doi.org/10.1016/j.jmateco.2008.12.002

Struik, D. (1953). *Lectures on analytic and projective geometry*. Addison-Wesley Publishing Company.

von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press.