STRICT DOMINANCE AND SYMMETRY

ALEXANDER R. PRUSS

Abstract. The Strict Dominance Principle that a wager always paying better than another is rationally preferable is one of the least controversial principles in decision theory. I shall show that (given the Axiom of Choice) there is a contradiction between Strict Dominance and plausible isomorphism or symmetry conditions, by showing how in several natural cases one can construct isomorphic wagers one of which strictly dominates the other. In particular, I will show that there is a pair of wagers on the outcomes of a uniform spinner which differ simply in where the zero degrees point of the spinner is defined to be but where one wager dominates the other. I shall also argue that someone who accepts Williamson's famous argument that the probability of an infinite sequence of heads is zero should accept the symmetry conditions, and thus has reason to weaken the Strict Dominance Principle, and I shall propose a restriction of the Principle to "implementable" wagers. Our main result also has implications for social choice principles.

1. Introduction

The Strict Dominance Principle is among the least controversial principles in decision theory: When one wager beats a second no matter what happens, the first wager certainly looks more rational than the second. Dominance is useful for arguments when one cannot assume that both sides of a discussion accept more controversial decision-theoretic principles, but where one can argue that some credential or decision-making process is irrational because it yields an outcome sure to be worse than another process's outcome. Thus, dominance has been used in pragmatic and non-pragmatic arguments for the rational requirement of probabilistic consistency in credences (e.g., [4], [8], [9]) and is of course central to discussions of the Newcomb Problem. Even Dutch Book arguments that show that some process leads to acceptance of a series of wagers that together result in a sure loss (e.g., [1], [13]) may just be a special case of dominance arguments, in that arguably what is objectionable about accepting a Dutch Book is that one is accepting a set of wagers that are collectively dominated by the status quo.

I shall argue, however, that there is a cost to the Strict Dominance Principle: someone who affirms it needs to either deny plausible symmetry principles or opt for a controversial narrowing of decision theory to what one might call "implementable wagers". Our main example of where symmetry and strict dominance conflict involves a paradoxical wager similar to the case of Vitali sets, whose existence is famously proved with the Axiom of

Choice, and hence our paradoxical wager will not be one where we have an explicit mathematical construction.

Suppose we have a sample space Ω of possible states. A wager W will be a real-valued function on Ω whose values represent utilities and which is bounded, i.e., there is a finite number M such that $|W(\omega)| < M$ for all states ω. The reason for the restriction to *bounded* wagers through most of the paper is to avoid paradoxes involving unbounded wagers, such as St. Petersburg.

There are several concepts of dominance. Given wagers W_1 and W_2, we say that:

(1) W_2 strictly dominates W_1 if and only if $W_2(\omega) > W_1(\omega)$ for all ω
(2) W_2 weakly dominates W_1 if and only if $W_2(\omega) \geq W_1(\omega)$ for all ω with the inequality being strict for at least one ω
(3) W_2 non-strictly dominates W_1 if and only if $W_2(\omega) \geq W_1(\omega)$ for all ω.

Additionally, if there is a credence function P from some algebra of subsets of Ω to $[0, 1]$, then we say:

(4) W_2 almost strictly dominates W_1 with respect to P if and only if $W_2(\omega) \geq W_1(\omega)$ for all ω and the inequality is strict except on some set $A \subseteq \Omega$ with $P(A) = 0$.[1]

Strict dominance entails almost strict dominance and, if P satisfies the axioms of finitely additive probability[2], then almost strict dominance entails weak dominance. These conditions naturally extend to unbounded wagers.

Write $W_1 \precapprox W_2$ to mean that W_2 is non-strictly preferred by the agent in question to W_1, and $W_1 \prec W_2$ provided that the agent strictly prefers W_2 to W_1, i.e., that $W_1 \precapprox W_2$ but not $W_2 \precapprox W_1$. Write $W_1 \approx W_2$ provided that the agent is indifferent between W_1 and W_2, namely $W_1 \precapprox W_2$ and $W_2 \precapprox W_1$. It is usual to assume that $\precapprox$ is a partial preorder, which is a reflexive and transitive relation, and often decision theorists assume preference structures that are total preorders, which have the additional property that at least one of $W_1 \precapprox W_2$ and $W_2 \precapprox W_1$ holds for any wagers W_1 and W_2. However, our main results will not need reflexivity, transitivity, or totality.

Say that $\precapprox$ satisfies the Principle of Strict (or Weak or Almost Strict, respectively) Dominance if for any wagers such that W_2 strictly (or weakly or almost strictly, respectively) dominates W_1, we have $W_1 \prec W_2$. And say that it satisfies the Principle of Non-Strict Dominance if whenever W_2 non-strictly dominates W_1, we have $W_1 \precapprox W_2$.

Because we did not assume reflexivity, it is formally possible for the preference structure to be empty—for the agent not to have any preference relations between any wagers. It is also formally possible for there to be

[1]This condition is related to the condition that $P(\{\omega : W_2(\omega) = W_1(\omega)\}) = 0$, but is weaker. For instance, it may be that the set $\{\omega : W_2(\omega) = W_1(\omega)\}$ is not P-measurable, even though it is contained in a set with P-measure zero.

[2]Actually, all we need is that $P(\Omega) \neq 0$.

some $\precsim$-unrelated wagers, ones that are not related by preference to any wagers, not even themselves. However, as soon as we have any of the four dominance principles defined in the previous paragraph, it will follow that for any wager W, we have $W-1 \precsim W \precsim W+1$, where $W+\alpha$ is the wager that pays α more than W no matter what, and hence no wager will be $\precsim$-unrelated. Each of the dominance principles thus implies that the preference structure extends to all wagers.

Now, certain sample spaces come along with intuitive symmetries. For instance, suppose that Ω is the set of outcomes of a uniform fair spinner, naturally identified with the set of points on the circumference of the unit circle. Then any rotation (and reflection, for that matter) is a plausible symmetry of Ω. We say that an agent's preference structure $\precsim$ is (strongly) invariant under a set G of permutations of Ω provided that we always have $W \approx W^g$, where W^g is the wager defined by $W^g(\omega) = W(g\omega)$ and g is in G. In the spinner case, if g is a rotation, then W^g is the wager you get by adding an extra rotation g to the end of a spin before calculating the payoff using W, and (strong) G-invariance says that adding an extra rotation doesn't affect the agent's preferences. Intuitively, if our spinner is known by a rational agent to be truly uniform, adding an extra rotation to the end of the spin shouldn't affect the agent's preferences. The wagers W and W^g are isomorphic in a very natural way—they just differ in where we put zero degrees on the circle when defining the payoffs.

Pruss [10] showed that the Principle of Weak Dominance is incompatible with rotational invariance. Fix any irrational number x. Let A is the set of points on the circle at $x, 2x, 3x, ...$ degrees, and suppose W is the wager that pays a dollar on A and zero outside it. Then if ρ is rotation by x degrees, W^ρ pays a dollar on the set A_0 of points at $0, x, 2x, ...$ degrees and zero outside A_0. Thus, W^ρ weakly dominates W, and so $\precsim$ cannot both satisfy rotational invariance and the Principle of Weak Dominance.

However, this is not a particularly impressive example. Because A is a countably infinite set, its Lebesgue measure is zero, and so it is reasonable to say that the probability of any non-zero payoff by either wager is zero on classical probability theory, or at best an infinitesimal if a non-classical theory is preferable. It does not seem particularly costly to say that one can ignore infinitesimally unlikely outcomes in one's decision theory, and be indifferent between wagers that differ in this way.

But I will show that, given the Axiom of Choice, for any irrational x and ρ defined just as above, there is a different wager V, obtained by tweaking a construction of Norton [6], such that V^ρ *strictly* dominates V. Then it is impossible to have a preference preorder $\precsim$ that satisfies both the Principle of Strict Dominance and rotational invariance. I will give the simple proof of this result in the next section, together with a simple general characterization of precisely when one can have both Strict Dominance and invariance under a group of symmetries, and some further examples. I will then discuss philosophical consequences for decision theory. In particular I will argue that we may well need to choose between Strict Dominance and the kind of symmetry reasoning that underlies Williamson's famous argument that the probability of an infinite sequence of heads is zero rather than a non-zero infinitesimal. One suggestion I will offer is that we should consider restricting our decision theory to "implementable" wagers. I will end with a brief sketch of some extensions, including to analogous results regarding Pareto conditions in social choice.

2. Main result and some examples

First, we prove the existence of our paradoxical wager in the case of the spinner. As before, let x be an irrational number and let ρ be rotation by x degrees. Define the relation $\sim$ on points on the unit circle S^1 by letting $a \sim b$ hold if and only if b can be obtained by rotating a by an integer multiple of x degrees about the center of the circle. It is easy to see that $\sim$ is reflexive, symmetric and transitive, and hence it divides the unit circle into a collection of equivalence classes. By the Axiom of Choice, let E be a set that contains exactly one element from each equivalence class. Given a point a on the circle, there is a unique integer n such that there exists a $b \in E$ such that a is the result of rotating b by nx degrees (here the irrationality of the number x is used).[3] Let $U(a) = n$. Observe that $U(\rho a) = U(a) + 1$ for any ρ.

Thus, U^ρ strictly dominates U. However, U is unbounded. To get a bounded wager, let $V(a) = \phi(U(a))$ for any strictly increasing bounded function ϕ on the reals (e.g., $\phi(y) = \arctan y$ or $\phi(y) = y/(1 + |y|)$). We will still have V^ρ strictly dominating V.

We can now generalize the above reasoning to show that *unless* a certain technical condition on the symmetries of Ω holds (condition (iv) below), then there is guaranteed to be a paradoxical wager like the above. The technical condition says that for every symmetry g, there is some point of Ω such that a finitely repeated application of g will return that point to itself. This condition was not met in our spinner case, because repeated rotation by x degrees will never return to the original point if x is irrational. On the other hand, if we *ad hoc* restricted our rotations to angles that are rational numbers of degrees, the technical condition would be met, since if g is rotation by p/q degrees where p and q are integers with $q \geq 1$, then $360q$ applications of g would return any point to itself. Moreover, it turns out that when the technical condition *is* satisfied, then we can prove that there is an invariant preference structure, and even one that is total.

[3]First, there is exactly one member b in E such that $a \sim b$, since E contains exactly one element from each equivalence class. Now suppose a can be obtained by rotating b by nx degrees as well as by mx degrees for integers n and m. We must show that $n = m$. But the only way rotation by nx degrees and one by mx degrees applied to b can both yield the same point a is if $nx - mx = (n - m)x$ is an integral multiple of 360. But since x is irrational, this can only happen if $n - m = 0$.

Proposition 1. *Assume the Axiom of Choice. Let G be a group of permutations of Ω. Then the following conditions are equivalent:*

- *(i) There is a G-invariant preference structure on wagers on Ω that satisfies the Principle of Strict Dominance.*
- *(ii) There is a G-invariant total preorder on wagers on Ω that satisfies the Principles of Strict Dominance and Non-Strict Dominance.*
- *(iii) There is no wager V and $g \in G$ such that V^g strictly dominates V.*
- *(iv) For every member g of G, there is a positive integer n and a member ω of Ω such that $g^n\omega = \omega$.*

Proposition 1 does not depend on any assumptions about probabilities (whether classical, risk-averse, qualitative, etc.), and while it is formulated above for bounded real-valued wagers, it works for wagers with any values that have a bounded subset of the same order type as the integers. It further has applications in completely non-probabilistic situations: see Section 4.[4]

The proof of the Proposition is given in the Appendix.

The simplest example where Strict Dominance is impossible is when Ω is a bidirectionally infinite fair lottery with tickets corresponding to the integers, and G consists of all forward/backwards shifts of ticket numbers. Thus, a member of G is a permutation ρ_n, for an integer n, such that $\rho_n m = n + m$. In this case, the Axiom of Choice is not needed, as we can exhibit the wager V explicitly as $V(n) = \phi(n)$ for any bounded strictly increasing function ϕ. But while this example is mathematically trivial, the possibility of infinite fair lotteries is quite philosophically controversial (see Pruss [12, Chapter 4] and Norton and Parker [7]).

Here is a less trivial case. Suppose first we have a bidirectionally (countably) infinite sequence of independent flips of a fair coin. For instance, we might have a line of people, stretching to infinity in both directions, and each one flips a fair coin, or we might have an infinite past and and an infinite future and each day one fair coin is flipped. We can represent a point ω in our space Ω as a function that assigns to each integer n either T or H, depending on whether the nth coin is tails or heads. We denote the set of functions from the set of integers $\mathbb{Z}$ to $\{T, H\}$ as $\Omega = \{T, H\}^{\mathbb{Z}}$.

In this example, one natural group G of symmetries are shifts by an integer amount. Thus, a permutation g in G shifts the results along the sequence, and is of the form ρ_m where ρ_m is a shift by m to the left: $(\rho_m\omega)(n) = \omega(m + n)$. In this example, condition (iv) is actually satisfied. For given any shift ρ_m, we have $\rho_m\omega = \omega$ if ω consists only of heads or only of tails. Thus, Proposition 1 tells us that there is a total preorder on Ω that satisfies the Principles of Strict and Non-Strict Dominance and is invariant under all shifts.

However, it is too soon for friends of dominance and symmetry to rejoice even in this case. For in this case we can show that there is no G-invariant preference that satisfies Almost Strict Dominance. To see this, say that a

---

[4] I am grateful to an anonymous reader for these observations.

heads-and-tails sequence ω in Ω is m-periodic provided that it repeats with a cycle of length m: i.e., $\omega(m+n) = \omega(n)$. And say that ω is periodic if for some finite m it is m-periodic. Let Ω_0 be the set of sequences in Ω that are *not* periodic. If g is a shift by a non-zero amount, then the only way we can have $g^n\omega = \omega$ for $n > 0$ is if ω is periodic. Thus, if replace Ω with Ω_0, then condition (iv) will fail. Thus, by the negation of (iii) there will be a paradoxical wager V on Ω_0 and a shift g such that V^g strictly dominates V.

Now observe that with respect to the classical probability measure P on Ω that treats all the coin tosses as fair and independent, the set of periodic sequences has zero probability. For, given any m, there are exactly 2^m sequences that are m-periodic, and hence there are only countably many periodic sequences.[5] But on the classical coin-flip probability measure, every individual sequence ω has probability zero, and hence by countable additivity, the set of all of them has probability zero. We can now take our paradoxical wager V on Ω_0 where V^g strictly dominates V, and extended it to a wager W on Ω by saying that $W(\omega) = V(\omega)$ for $\omega \in \Omega_0$ and $W(\omega) = 0$ for a periodic ω. Then $W^g(\omega) > W(\omega)$ for every non-periodic ω, and $W^g(\omega) = W(\omega)$ for every periodic ω, so W^g almost strictly dominates W. It follows that there is no G-invariant preference $\precsim$ for wagers on Ω that satisfies Almost Strict Dominance, since if $\precsim$ satisfied Almost Strict Dominance, we would have $W \prec W^g$, contrary to G-invariance.

3. Philosophical consequences

Proposition 1 shows that the following three theses are contradictory:

(5) The Axiom of Choice holds.
(6) Ideal agents' preferences always satisfy Strict Dominance.
(7) Ideal agents' preferences are sometimes invariant under symmetries that fail condition (iv) of Proposition 1, such as rotational symmetry for uniform spinners or translational symmetry for infinite countable fair lotteries.

I will take the Axiom of Choice for granted—it is widely accepted by mathematicians, and a serious discussion would move us from decision theory to the philosophy of mathematics. Thus, we have a choice between rejecting Strict Dominance and rejecting the relevant cases of symmetry.

Furthermore, the following thesis is very plausible:

(8) If ideal agents' preferences always satisfy Strict Dominance, they also always satisfy Almost Strict Dominance.

In addition to the strong intuitive plausibility of (8), we can argue for it as follows. Suppose (8) is not true, so that ideal agents' preferences always satisfy Strict Dominance but not always Almost Strict Dominance. Thus we will have a case where an ideal agent, call her Alice, has a credence function P and two wagers W_1 and W_2 have the property that W_2 almost strictly

[5] The union of a countable sequence of finite sets is countable.

dominates W_1 with respect to P, but Alice does not prefer W_2 to W_1. Let A be an event with $P(A) = 0$ such that $W_2 > W_1$ everywhere outside of A. Plausibly, if Alice were to gain complete "Cartesian" certainty that A did not occur, she would prefer W_2 to W_1: events one is completely certain did not occur not be included in the relevant sample space, and with the space restricted to the outside of A, there would be strict dominance between the wagers W_2 and W_1. Now, it is natural to treat probability one as a kind of certainty, and so Alice may well have a kind of certainty that A did not occur, but it need not be a complete certainty (e.g., even if the probability of an infinite sequence of heads on fair independent tosses is zero, one does not have complete "Cartesian" certainty that such a sequence won't transpire). The difference between complete certainty and probability one is presumably behind why Alice, despite being an ideal agent and ideal agents' preferences satisfying Strict Dominance, does not strictly prefer W_2 to W_1.

Still, when a wager W_2 is strictly preferred to W_1 outside of an event A, it seems that the only reasons an ideal agent might have *not* to prefer W_2 to W_1 would either be (a) a worry that W_1 is preferable to W_2 on A or (b) a worry that the case where A does not occur is probabilistically negligible or (c) skepticism about the Principle of Strict Dominance. Now we've assumed that ideal agents do respect the Strict Dominance Principle. Furthermore, Alice should not worry that W_1 is preferable to W_2 on A, since even on A we have $W_2 \geq W_1$, and it is the case where A occurs, rather than the case where it does not, that is probabilistically negligible. Hence none of (a)–(c) apply, and Alice should prefer W_2 to W_1.

If we accept (8), the stakes in the conflict between Strict Dominance and symmetry are raised. For even if one is skeptical of countably infinite fair lotteries (e.g., [6] or [12]) and of completely rotationally symmetric spinners, it is very plausible that preferences between wagers concerning a bidirectionally infinite sequence of independent fair coin flips should be invariant under translations.

In particular, if we accept Almost Strict Dominance and (8), we will now have a new reason to be skeptical of Williamson's [17] famous argument that an infinite sequence of heads has probability zero rather than being a non-zero infinitesimal. Williamson's argument supposes a unidirectionally infinite sequence of independent fair coins being tossed, which we can imagine arranged at spatial or temporal locations $1, 2, 3, ...$ Let the probability of that sequence being all heads be ε. Williamson then has us suppose another independent fair coin being tossed at location 0. The probability of this larger sequence of coins being all heads then equals $(1/2)e$. But the two sequences are isomorphic in a sense relevant to probabilistic reasoning, so $\varepsilon = (1/2)\varepsilon$, which is only possible if $\varepsilon = 0$.

Now, if we accept the isomorphism step in Williamson's argument, then we should likewise hold that a *bidirectionally* infinite sequence of independent fair coin tosses is epistemically isomorphic to a sequence translated by one coin to the left. The intuitions here are the same. But an ideal agent would not have a strict preference between epistemically isomorphic wagers. And yet Almost Strict Dominance forces such a preference, *if* translations result in epistemically isomorphic situations. Thus, if we accept Almost Strict Dominance, we need to reject that translations of coin flips result in epistemically isomorphic situations, and hence we should reject Williamson's argument. Those readers who find Williamson's argument compelling will thus have reason to reject the Principle of Almost Strict Dominance, and, given (8), the Principle of Strict Dominance.

Of course, when we reject a plausible thesis, like Strict Dominance, we should try to find another in its place. One might consider replacing Strict Dominance with Non-Strict Dominance, as Non-Strict Dominance can be satisfied no matter what symmetries we want. However, the Principle of Non-Strict Dominance is quite a weak thesis, given that it can be satisfied trivially by the total preference preorder where one is indifferent between all pairs of wagers.

But there is a more interesting way to weaken Strict Dominance. As defined, the Principle applies to all wagers, and wagers are just bounded mathematical functions. When preference structures were introduced, it was noted that each of the four dominance principles (Strict, Weak, Non-Strict or Almost Strict) was sufficient to imply that every wager was preference-related to some other wager. But not every mathematical function can be the payoff function for a physically implementable game. For instance, in the circle case, the choice set E containing one element from each $\sim$-equivalence class is a non-measurable set similar to a Vitali set. Norton [6] considers an infinite lottery very similar to one obtained by spinning our spinner and applying our function U to the outcome, and argues that such a lottery is not implementable because there is no way to check whether the spinner lands in a particular Vitali set. The same point applies to our case. We can thus weaken our dominance principles to apply only to physically implementable wagers. This would leave open the possibility that some wagers (say, some or all the unimplementable ones) are not preference-related to any wagers or are not preference-related to any wagers distinct from themselves.

It is worth noting that in the circle case, a wager involving the paradoxical property that W^ρ strictly dominates W for some rotation ρ is *always* a Lebesgue non-measurable function. For if W^ρ strictly dominated W and yet W was measurable and bounded, we would absurdly have:

$$0 = \int_{S^1} W^\rho \, d\lambda - \int_{S^1} W \, d\lambda = \int_{S^1} (W^\rho - W) \, d\lambda > 0,$$

where λ is Lebesgue measure, the first equality follows from the rotational invariance of Lebesgue integrals, the second from their linearity, and the final inequality follows from strict dominance and the fact that the Lebesgue integral of a strictly positive function is strictly positive. A similar argument shows that in the case of the bidirectionally infinite sequence of coin tosses, any wager W such that W^ρ almost strictly dominates W for some translation ρ will also fail to be measurable with respect to the standard product measure on $\{H,T\}^{\mathbb{Z}}$.[6] And it is plausible that non-measurable functions are not implementable as actual games.[7]

Formally, we then have this fact contrasting with Proposition 1:

Proposition 2. *Suppose μ is a countably additive probability measure on a σ-algebra $\mathcal{F}$ of measurable subsets of Ω that is invariant under a group G, i.e., if ρ in G and $A \in \mathcal{F}$, then $\rho A \in \mathcal{F}$ and $\mu(\rho A) = \mu(A)$. Define the preference relation $W_1 \precsim W_2$ if and only if W_1 and W_2 are μ-measurable and $\int_\Omega W_1\,d\mu \le \int_\Omega W_2\,d\mu$. Then $\precsim$ is transitive, is reflexive on the measurable wagers, satisfies the restriction of Strict Dominance Principle to measurable wagers, and has the symmetry property that $W \approx W^\rho$ for any measurable wager W.*

This follows from the fact that the integral of W and W^ρ is the same if the measure is G-invariant and $\rho \in G$, and that if $W_1 < W_2$ everywhere, then $\int_\Omega W_1\,d\mu < \int_\Omega W_2\,d\mu$ for any measure μ. Lebesgue measure on the circle will be invariant under rotations, and in many other contexts—like the coin toss one—there will be other invariant countably additive measures.[8]

Restricting Strict Dominance to physically implementable wagers thus may allow us to preserve symmetry, assuming only measurable functions are implementable. Perhaps, though, we have higher ambitions for decision theory, and think general principles such as Strict Dominance should apply at least to all metaphysically possible situations. However, not only is it unlikely that a game involving a non-measurable payoff function is physically implementable, it may even be metaphysically impossible. The only way we know how to knowingly implement selections that use the Axiom of

[6]The only difference is that in the last step of the proof where we will need to show that $\int_{\{H,T\}^{\mathbb{Z}}} (W^\rho - W)\,dP > 0$, we won't have a guarantee that $W^\rho - W > 0$ everywhere, but only that this happens outside of a set of measure zero. But that's enough for the integral to be strictly positive.

[7]Solovay [14] has famously shown that assuming a certain large cardinal assumption, the existence of a non-measurable set requires some version of the Axiom of Choice. If the assumption holds, then any paradoxical wager like the one in the spinner case requires the Axiom of Choice. Note, however, that the fact that the proof of something requires the Axiom of Choice does not prove that the thing cannot be explicitly constructed, as Kanovei and Shelah's [5] construction of a free ultrafilter on an infinite set shows.

[8]Note that Proposition 2 becomes false if "countably additive" is replaced by "finitely additive", at least given the Axiom of Choice. Let $G = \Omega = \mathbb{Z}$ be the integers, acting on themselves by addition. Since $(\mathbb{Z},+)$ is a commutative group, it is amenable [16, Theorem 12.4], i.e., there is a finitely additive $\mathbb{Z}$-invariant probability measure μ on $\mathbb{Z}$. Invariance implies that every singleton has the same measure, and finite additivity forces that measure to be zero, so again by finite additivity the measure of every finite set is zero. Let W_1 be the wager that is zero everywhere and let $W_2(n) = 1/(1+|n|)$. Then W_2 strictly dominates W_1, but $0 \le \int_\Omega W_2\,d\mu \le 1/n$ for every positive integer n, since $W_2 \le 1/n$ everywhere except on the finite set $\{-(n-2),\dots,n-2\}$ which has μ-measure zero, so $\int_\Omega W_2\,d\mu = 0$, which is also what the integral of W_1 equals.

Choice involves nested supertasks and it is far from clear that supertasks are metaphysically possible [12]. Indeed, we might even take the tension between Strict Dominance and symmetry, together with the apparent implementability of the Choice-selections with supertasks, as a further argument against the metaphysical possibility of supertasks.

We might, on the other hand, choose to weaken the symmetry condition. Again, one way is to apply it only to physically or maybe even metaphysically implementable wagers.

Another approach is to weaken strong G-invariance to weak G-invariance, where instead of requiring that $W \approx W^g$, we require that for all W_1, W_2 and g, we have $W_1 \precapprox W_2$ if and only if $W_1^g \precapprox W_2^g$. Things look much better for weak G-invariance. For instance, if we define a preference preorder by stipulating $W_1 \precapprox W_2$ just in case W_2 non-strictly dominates W_1, then we satisfy both Strict Dominance and weak G-invariance for any group of permutations G. Moreover, it follows from [11, Theorem 2] that whenever G is a commutative group (this will be true for the rotations in our spinner example and the shifts in our bidirectionally infinite lottery and coin toss examples), this partial preorder $\precapprox$ can be extended to total preorder.

However, weak G-invariance does not capture our symmetry intuitions unless it implies strong G-invariance. In the case of the spinner, it is not only intuitive that our relative rankings between a pair of wagers should be unchanged by rotating one of the wagers, but also that they should be unchanged by rotating *one* of the wagers, and that condition immediately implies strong G-invariance (just apply it in the case where the two wagers are the same).

Furthermore, in the case of the spinner, the countably infinite fair lottery and the bidirectionally infinite sequence of coin tosses, it is just as intuitive to think that we should have weak invariance under reflections[9] as that we should have weak invariance under rotations or translations. However, weak invariance of a total preorder under all reflections implies strong invariance under all reflections by [11, Prop. 2], and any rotation (in the spinner case) or translation (in the lottery and coin cases) can be generated by a pair of reflections.[10] Thus, at least in the case of a totally preordered preference structure, weak invariance under reflections implies strong invariance under rotations or translations (depending on the case), and we have seen that strong invariance cannot be had.

---

[9]In the case of the spinner, we can reflect the result in a line through the center of the spinner. In the lottery, for any fixed m such that $2m$ is an integer, we can map an outcome n to its reflection $m - n$ around m, and in the coin toss example, for any such m, we can map ω to ω_m defined by $\omega_m(n) = \omega(m - n)$.

[10]A rotation by an angle α can be generated by reflecting about a line at angle $\alpha/2$ and then a line at angle $3\alpha/2$; a translation by a distance x can be generating by reflecting about the point $x/2$ and then the point $3x/2$.

We thus have a choice: To limit the scope of decision to implementable wagers and argue that the paradoxical wagers in this paper are not implementable, to reject Strict Dominance, or to reject very plausible symmetry principles, including the one underlying the Williamson argument.

4. Some extensions

While the vocabulary of wagers has been used so far, and Ω in Proposition 1 has been interpreted as a sample space, the Proposition as a mathematical result has implications that go beyond these interpretations.[11]

For instance, we can reinterpret Ω as a set of sites of value (e.g., persons at spatiotemporal locations) and a wager as a deterministic distribution of goods across sites, and get a result for non-probabilistic social choice principles. For the Proposition implies that a preference structure on distributions of goods (even ones of bounded value) across the sites that is invariant under some group G of symmetries of the sites can only satisfy the weak Pareto condition that a distribution that is better for everyone is always preferred to a distribution that is worse for everyone if the symmetries satisfy condition (iv) of the Proposition (i.e., if there is a symmetry g that does not move any point of Ω around a finite cycle). For instance, if the persons are located at all the intersection points of a two- or three-dimensional rectangular grid, and the symmetries are translations along the axes of the grid, then condition (iv) fails, and no symmetry-invariant preference structure satisfies the weak Pareto condition.[12] If one is confident of the weak Pareto condition, this provides another argument that spatiotemporal location may be morally significant (cf. [3]), or it might make one suspicious of the weak Pareto condition.

Or, more generally, following Easwaran [2] we can consider Ω to be a set $S \times \Omega_1$ of pairs (z, ω) where $z \in S$ is a site of value and $\omega \in \Omega_1$ is a location in a probabilistic sample space, allowing one to consider probabilistic situations where values are not aggregated into a single value. In this case, G can be a set of symmetries that act on the sites, on the sample space Ω_1, or on combinations of the two, and once again we can conclude that if condition (iv) fails, then any symmetry-invariant preference structure fails the weak Pareto condition that if $W_1 < W_2$ everywhere on Ω, then W_2 is strictly preferable.

Finally, it is worth noting that if we do not require the values to be bounded, then the proof of the Proposition shows that we can replace the

---

[11] I am grateful to an anonymous reader for pointing me to these applications.

[12] That said, in this special case the non-existence of such a preference structure does not need the Axiom of Choice and is in fact obvious. Suppose the sites have integer coordinates (x, y, z). Let $W(x, y, z) = \phi(x)$ for any strictly increasing bounded function ϕ (e.g., arctan) and let τ be translation by one unit to the right along the x-axis. Then $W(x, y, z) < W(x + 1, y, z) = W^\tau(x, y, z)$ for all sites (x, y, z), and hence a translation-invariant preference structure cannot satisfy the weak Pareto condition.

Strict Domination Principle with what we might call the Very Strict Domination Principle that if $W_1 + \varepsilon < W_2$ everywhere on Ω for some fixed $\varepsilon > 0$, then W_2 is strictly preferable to W_1. In the social choice setting, this is yields a result for an even weaker Pareto condition, also discussed by Easwaran [2, Section III].[13]

Appendix: Some proofs

Proof of Proposition 1. We will show (i)→(iii)→(iv)→(ii)→(i).

That (i) implies (iii) is clear: if we had V^g strictly dominating V and yet $\precapprox$ was a G-invariant preference structure satisfying Strict Dominance, then by Strict Dominance we would have $V \prec V^g$, contrary to G-invariance.

Now we show that (iii) implies (iv). Assume (iv) is false. Thus there is a symmetry g such that $g^n\omega \neq \omega$ whenever $n > 0$ and $\omega \in \Omega$. More generally, it follows that if $g^n\omega = g^m\omega$, then $n = m$. For, otherwise, we could suppose without loss of generality that $n > m$ and we would have $g^{n-m}\omega = \omega$.

Let $a \sim b$ if and only if $a = g^n b$ for some integer n. Then $\sim$ is an equivalence relation, and we can let E contain exactly one element from each equivalence class by Choice. Given $a \in \Omega$, let $b \in E$ be such that $a = g^n b$ for some integer n. The integer n is unique, since otherwise we would have $g^n b = g^m b$ for distinct n and m. Let $U(a) = n$, observe that $U^g(a) = U(a) + 1$ and, as before, let $V(a) = \phi(U(a))$ for a bounded strictly increasing ϕ. Then V^g strictly dominates V, and we have not-(iii).

Now we show (iv) implies (ii). Assume (iv). Say that $V \precapprox W$ just in case there is a $g \in G$ such that $V(\omega) \leq W^g(\omega)$ for all $\omega \in \Omega$. Then $\precapprox$ is a G-invariant partial preorder, and it clearly satisfies the Principle of Non-Strict Dominance. We now show that it satisfies Strict Dominance as well. For suppose it does not, so that there are W_1 and W_2 such that $W_1 < W_2$ everywhere but not $W_1 \prec W_2$. By Non-Strict Dominance, we have $W_1 \precapprox W_2$. Thus, for $W_1 \prec W_2$ to fail, we must also have $W_2 \precapprox W_1$. Hence there is a $g \in G$ such that $W_2 \leq W_1^g$ everywhere. Therefore, $W_2 < W_2^g$ everywhere, since $W_1 < W_2$ everywhere. By (iv), there is a positive n and an ω such that $g^n\omega = \omega$. We then have $W_2(\omega) < W_2(g\omega) < W_2(g^2\omega) < \cdots < W_2(g^n\omega)$, which contradicts $g^n\omega = \omega$.

However, $\precapprox$ may be only a partial preorder, and we need a total one. To that end, let $[W]$ be the equivalence class of the wager W under the relation $\approx$ (where, recall, $W_1 \approx W_2$ if and only if $W_1 \precapprox W_2$ and $W_2 \precapprox W_1$). Define the partial order $\preceq$ on these equivalence classes by stipulating that $[W] \preceq [V]$ if and only if $W \precapprox V$ (this is well defined because $\precapprox$ is transitive). By the Szpilrajn order extension theorem [15] (this uses the Axiom of Choice), we can extend $\preceq$ to a total order $\preceq^*$. Now define $W_1 \precapprox^* W_2$ if and only if $[W_1] \preceq^* [W_2]$. Then $\precapprox^*$ is easily seen to be a total preorder that extends $\precapprox$. It is strongly G-invariant because it extends the strongly G-invariant

[13] I am grateful to [anonymized] for [anonymized], and to two anonymous readers for a careful reading and a number of suggestions that have significantly improved the paper.

preorder $\precapprox$. It remains to check that $\precapprox^*$ satisfies the Principles of Non-Strict and Strict Dominance. Non-Strict Dominance follows from the fact that $\precapprox^*$ extends $\precapprox$ and the latter satisfies Non-Strict Dominance. That leaves the Strict case. Suppose W_2 strictly dominates W_1. By Strict Dominance for $\precapprox$, we have $W_1 \prec W_2$. Therefore, $[W_1] \preceq [W_2]$ and not $[W_2] \preceq [W_1]$. Hence $[W_1] \preceq^* [W_2]$, since $\preceq^*$ extends $\preceq$. Since we do not have $[W_2] \preceq [W_1]$, we have $[W_2] \neq [W_1]$, and since $\preceq^*$ is an order, and not merely a preorder, it follows that we do not have $[W_2] \preceq^* [W_1]$. Thus, we have $W_1 \precapprox^* W_2$ but not $W_2 \precapprox^* W_1$.

Thus, (iv) implies (ii).

Finally, (ii) trivially implies (i), which completes the proof. □

References

[1] de Finetti, Bruno. 1937. "Foresight: Its logical laws, its subjective sources." In Henry E. Kyburg and Howard E.K Smokler (eds.), *Studies in Subjective Probability*, Huntington, NY: Kreiger Publishing.

[2] Easwaran, Kenny. 2021. "A new method of value aggregation" *Proceedings of the Aristotelian Society* 121:299–326.

[3] Jonsson, Adam and Peterson, Martin. 2020. "Consequentialism in infinite worlds", *Analysis* 80:240–248.

[4] Joyce, J. M. 1998. "A nonpragmatic vindication of probabilism", *Philosophy of Science* 65:575–603.

[5] Kanovei, V. and Shelah, S. 2004. "A definable nonstandard model of the reals". *Journal of Symbolic Logic* 69:159–164.

[6] Norton, John D. 2020. "How NOT to build an infinite lottery machine", *Studies in History and Philosophy of Science* 82:1–8.

[7] Norton, John D., and Parker, Matthew W. Forthcoming. "An Infinite Lottery Paradox", *Axiomathes*.

[8] Pettigrew, Richard. 2016. *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.

[9] Pettigrew, Richard. 2021. "On the expected utility objection to the Dutch Book argument for probabilism", *Noûs* 55:23–38.

[10] Pruss, Alexander R. 2013. "Null probability, dominance and rotation", *Analysis* 73:682–685.

[11] Pruss, Alexander R. 2014. "Linear extensions of orders invariant under abelian group actions", *Colloquium Mathematicum* 137:117–125.

[12] Pruss, Alexander R. 2018. *Infinity, Causation and Paradox*, Oxford: Oxford University Press.

[13] Ramsey, F. P. 1931. "Truth and probability." In: *The Foundations of Mathematics and Other Lgoical Essays*, Routledge and Kegan Paul, 156–198.

[14] Solovay, R. M. 1970. A Model of Set-Theory in Which Every Set of Reals is Lebesgue Measurable, *Transactions of the American Mathematical Society* 92/1: 1–56

[15] Szpilrajn, Edward. 1930. "Sur l'extension de l'ordre partiel", *Fundamenta Mathematicae* 16:386–389.

[16] Tomkowicz, G., and Wagon, S. 2016. *The Banach Tarski Paradox*, 2nd ed. Cambridge University Press: Cambridge.

[17] Williamson, Timothy. 2007. "How probable is an infinite sequence of heads?", *Analysis* 67:173–180.