Parity, Prospects, and Predominance

Ryan Doody

1 Introduction

Suppose I offer you a choice between an all-expense-paid Alpine ski vacation (A) and an all-expense-paid beach vacation (B). There are pros and cons to each. You might ultimately prefer one to the other. Or you might be completely indifferent. Or — if you're anything like me — you might feel torn. Suppose that, after considering all of the various pros and cons, you fail to have any of the three traditional preference-relations toward the vacations: you don't strictly prefer A to B, you don't strictly prefer B to A, and you aren't completely indifferent between them either. To make it clear that you aren't indifferent, let's suppose that you also don't prefer the Alpine ski vacation plus a dollar (A^+) to the beach vacation (and *vice versa*) and let's suppose that you don't prefer the beach vacation plus a dollar (B^+) to the Alpine ski vacation. Your preferences are *incomplete* (they violate the Completeness Axiom) — not because you've been careless, or lazy, or irrational but rather because you take the values of these prizes to be incommensurable. Let's, following [Chang, 2002], say that you regard A and B as *on a par*.

What does rationality require of you when facing a choice between options that you regard as on a par? Here's a plausible thought: If you regard A and B as on a par, you are not rationally required to choose A and you are not rationally required to choose B. That seems right enough to me. This paper is concerned with a related, but more difficult question: What does rationality require of you when facing a choice between *risky* options whose outcomes you regard as on par? Standard decision theory is silent on this issue.

Hare [2010] develops two different ways of generalizing decision theory to cases in which you regard the potential outcomes of your options to be on a par. He cautiously sides with one of these proposals, which he calls *Prospectism*, over the other. (Roughly, *Prospectism* says that you should perform the option with the best prospects, where an option's prospects is a probability-distribution over its potential outcomes). It's been pointed out that *Prospectism* violates a "dominance" principle: it will sometimes require you to choose an option that, no matter how the world is, cannot be better than the alternatives.[1] But the arguments given against violating a dominance principle like this are too quick. I offer a response on *Prospectism*'s behalf.

I will go on to argue, however, that things are worse for *Prospectism*: it violates a weaker, and better motivated, dominance principle. Sometimes, *Prospectism* requires you to choose an option that, no matter how the world is, cannot be better than the alternatives and is such that there are some ways the world might be in which it is worse. In cases like these, you are in a position to know of one of your options that it's not worse than the alternatives and that it might (or even likely!) be better. This, I contend, provides you with a compelling reason in favor of choosing it. And yet *Prospectism* says that it would be irrational for you to do so. This strikes me as implausible. And, as I'll argue, it undermines one of the strongest considerations in support of *Prospectism*: that we're required to do what we have the most reason to do.

2 *Prospectism* and Opaque Sweetening

Consider the following case, from Hare [2010]:

> **Vacation Boxes.** There are two opaque boxes: a Larger box (L) and a Regular box (R). A fair coin has been tossed. If the coin landed heads, then a voucher for an all-expenses-paid Alpine ski vacation (A) was

[1] I've put 'dominance' in scare quotes because, unlike the more traditional dominance-relations familiar from standard decision and game theory, the notion here isn't asymmetric: as we'll see, it's possible for two options to "dominate" each other (and every option trivially "dominates" itself). This might strike you as a misuse of the term. Hence, the scare quotes.

placed in the Larger box and a voucher for an all-expenses-paid beach vacation (B) was placed in the Regular box; if the coin landed tails, then B was placed in the Larger box and A was placed in the Regular box. In either case, you don't know which prize is in which box.

$$\textbf{Larger box} = \begin{cases} A & \text{if Heads} \\ B & \text{if Tails.} \end{cases} \qquad \textbf{Regular box} = \begin{cases} B & \text{if Heads} \\ A & \text{if Tails.} \end{cases}$$

Now imagine that \$1 is added to the Larger box. If you choose the Larger box, you will win whichever prize it contains plus a \$1. Nothing is added to the Regular box. You are asked to choose one of the two boxes, taking home whichever prize is in the box you choose.

	Heads	Tails
Take Larger box	A^+	B^+
Take Regular box	B	A

Does rationality require you to take the Larger box, or is it rationally permissible to take either?

Standard expected utility theory says nothing about cases like these because in order for *expected* utility to be well-defined, *utility* must be well-defined. But if you have incomplete preferences (as you do here), your preferences cannot be represented with a *single* utility-function.[2]

Hare's *Prospectism* is one way of generalizing expected utility theory to cases in which your preferences fail to be complete. There are two steps. First, we represent your preference-like attitudes with the *set* of utility-functions characterizing all of the coherent extensions of your preference ordering. An ordering, $\succeq^+$, is a *coherent extension* of a partial ordering, $\succeq$, just in case:

[2] Here is why you are unable to place a single, absolute value on any of these outcomes. Suppose, to the contrary, that you could. You assign the number $r \in \mathbb{R}$ to A: $u(A) = r$. Because you don't prefer A to B, the number you assign to B, $u(B)$, cannot be less than r. Because you don't prefer B to A, $u(B)$ also cannot be greater than r. Therefore, $u(B) = r$. And, because you prefer A^+ to A, it must be that $u(A^+) > r$. But, because you don't prefer A^+ to B, it cannot be the case that $u(A^+) > r$. And that's a contradiction.

(i) $\succeq^+$ is complete, and

(ii) for any outcomes X, Y, $X \succeq^+ Y$ if $X \succeq Y$.

So, for example, if your incomplete preference ordering ranks outcome X ahead of outcome Y, then *every* complete preference ordering included in the set will likewise rank X ahead of Y; and so on and so forth; if, however, outcome X and outcome Y don't stand in any of the three traditional preference-relations, then, for each of the ways the two can be ranked, there will be a complete preference ordering included in the set that does rank them that way.[3]

Second, we then apply the traditional machinery of expected utility theory to each of the complete orderings in your set. You should prefer one option to another just in case every complete ordering in your set ranks things that way; you should be indifferent between two options just in case every complete ordering in your set ranks them that way; etc.[4]

---

[3] Every partial ordering can be represented, in the manner described, by a set of complete orderings. The converse, however, doesn't hold: there are sets of complete orderings that cannot be faithfully represented by a partial ordering. Here's an example. Suppose you are deciding between three dessert options: an apple pie (A), a bowl of blueberries (B), and a cantaloupe cake (C). And, at least as far as desserts are concerned, you only care about two things: how *healthy* the dessert is, and how *delicious* it is. Suppose that A is the most delicious, B is the least delicious, and C is just slightly more delicious than B; and suppose that B is the healthiest option, A is the least healthy option, and C is just slightly healthier than A. Consequently, in terms of your all-things-considered preferences, none of the three options stand in any of the traditional preference-relations to any of the others. But, in such a case, we might want to represent your motivational-state with a set of complete orderings which includes orderings that rank C ahead of A and C ahead of B, but *doesn't* include any orderings that rank C ahead of *both* A and B. In other words, there are no admissible ways of evaluating your options, resolving your concern for health and your concern for deliciousness, according to which C is the dessert that is most desirable to you. (See [Levi, 1985, 2008] for a discussion of cases with this structure.) This distinction won't matter for our purposes, however.

[4] In addition to Hare's *Prospectism*, there are a number of views that have a very similar structure. See for example: I.J. Good's *Quantizationism* [Good, 1952]; Isaac Levi's *V-admissibility* [Levi, 1986, 2008]; Amartya Sen's *Intersection Maximization* [Sen, 2004]; and [Weirich, 2004]. Among economists, views of this general nature are nearly the only game in town. See, for example, [Dubra et al., 2004], [Evren and Ok, 2011], [Galaabaatar and Karni, 2013], [Ok et al., 2012]. Some of these views differ from the others in some important respects. But these differences won't matter for our purposes because each of the views recommend taking the Larger box over the Regular one. There are also a number of decision theories designed to handle similar cases that arise not because of incomplete preferences but because of imprecise (or unsharp) credences: for example, Susanna Rinard's *Moderate* [Rinard, 2015]; Weatherson's *Caprice* [Weatherson, 2008]; and [Joyce, 2010].

In Vacation Boxes, *Prospectism* entails that you rationally ought to take the Larger box; if you take the Regular box, then, according to *Prospectism*, you've behaved irrationally.

Here's why. You prefer A^+ to A and B^+ to B, so every utility-function in your set ranks A^+ ahead of A and ranks B^+ ahead of B. Let u^* be an arbitrary utility-function in your set.

$$u^*(\textit{take Larger}) = \frac{1}{2} \cdot u^*\left(A^+\right) + \frac{1}{2} \cdot u^*\left(B^+\right)$$
$$u^*(\textit{take Regular}) = \frac{1}{2} \cdot u^*\left(B\right) + \frac{1}{2} \cdot u^*\left(A\right)$$

No matter how u^* ranks A vs B, $u^*(\textit{take Larger}) > u^*(\textit{take Regular})$.[5] And because u^* was chosen arbitrarily, it follows that every utility-function in your representor ranks taking the Larger box ahead of taking the Regular box. So, you ought to prefer taking the Larger box to taking the Regular box.

3 Dominance Reasoning

As several philosophers have pointed out, *Prospectism* conflicts with dominance reasoning. (See, for example, [Bales et al., 2014], [Hare, 2010], [Schoenfield, 2014], and [Rabinowicz, 2016]). There is some sense in which taking the Regular box "dominates" taking the Larger box: no matter how the coin has landed, you won't prefer the outcome that would result from taking the Larger box to the outcome that would result from taking the Regular box. Taking the Regular box can do no worse than taking the Larger box. Does this, then, render taking the Regular box rationally permissible?

Perhaps, but not obviously so. We're use to dominance reasoning in contexts in which preferences are complete and value-relations are trichotomous, so it's worth

[5] Here's why. $u^*(\textit{take Larger}) > u^*(\textit{take Regular})$ iff $u^*(\textit{take Larger}) - u^*(\textit{take Regular}) > 0$. Because every utility-function in your set ranks A^+ ahead of A and ranks B^+ ahead of B, $u^*\left(A^+\right) > u^*\left(A\right)$ and $u^*\left(B^+\right) > u^*\left(B\right)$.

So, $u^*\left(A^+\right) - u^*\left(A\right) + u^*\left(B^+\right) - u^*\left(B\right) > 0$.

So, $\frac{1}{2}\left(u^*\left(A^+\right) - u^*\left(A\right)\right) + \frac{1}{2}\left(u^*\left(B^+\right) - u^*\left(B\right)\right) > 0$.

Thus, $\frac{1}{2}\left(u^*\left(A^+\right) + u^*\left(B^+\right)\right) - \frac{1}{2}\left(u^*\left(B\right) + u^*\left(A\right)\right) > 0$.

being extra careful in cases, like this one, in which parity is involved. For the sake of diligence, then, allow me to offer a taxonomy of "dominance"-relations.

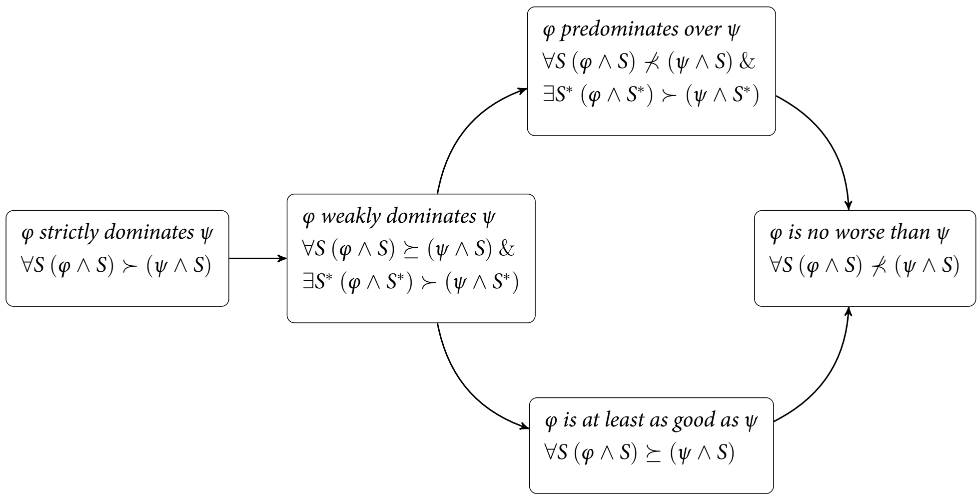

Figure 1: Logical Strengthen of Dominance Relations

Assuming that your preferences are complete, we can distinguish between three different ways one option might "dominate" another: an option can *strictly dominate* another, or it can *weakly dominate* another, or it can be *at least as good as* another.

Let $\mathbf{S} = \{S_1, S_2, \ldots, S_n\}$ be a partition of the ways the world might be. Option φ *strictly dominates* option ψ when, for every $S \in \mathbf{S}$, you prefer $(\varphi \wedge S)$ to $(\psi \wedge S)$. Informally, we'll say that one option strictly dominates another just in case it always does better. Option φ *weakly dominates* ψ when, for every $S \in \mathbf{S}$, you weakly prefer $(\varphi \wedge S)$ to $(\psi \wedge S)$ — that is, you either prefer the former outcome to the latter *or* you are indifferent between the two — and there is some $S^* \in \mathbf{S}$ such that you prefer $(\varphi \wedge S^*)$ to $(\psi \wedge S^*)$. Informally, we'll say that one option weakly dominates another just in case it always does at least as well and sometimes does better. Lastly, let's say that an option φ *is at least as good as* ψ when, for every $S \in \mathbf{S}$, you weakly prefer $(\varphi \wedge S)$ to $(\psi \wedge S)$. Informally, one option is at least as good as another just in case it always does at least as well: that is, no matter how the world turns out to be, you either prefer it or are indifferent.

If we allow for parity, there are two more notions. Let's say that an option φ *predominates over* ψ when, for every $S \in \mathbf{S}$, you don't prefer $(\psi \wedge S)$ to $(\varphi \wedge S)$ and there is some $S^* \in \mathbf{S}$ such that you prefer $(\varphi \wedge S^*)$ to $(\psi \wedge S^*)$. Informally, one option predominates over another if it never does worse and sometimes does better. Predominance is like weak dominance but weaker. And, finally, let's say that an option φ *is no worse than* ψ when, for every $S \in \mathbf{S}$, you don't prefer $(\psi \wedge S)$ to $(\varphi \wedge S)$. Informally, one option is no worse than another just in case it never does worse.

Strict dominance is the strongest of these notions. It entails the other four. Weak dominance is the next strongest. It entails the remaining three. If φ predominates over ψ, it follows that φ is no worse than ψ. But it doesn't follow that φ is at least as good as ψ. If φ is at least as good as ψ, it follows that φ is no worse than ψ. But it doesn't follow that φ predominates over ψ.

These five notions are potentially relevant to questions about what rationality requires because, when they hold between two options, you're in a position to know something about how the actual values of those options compare. For example, if φ strictly dominates ψ, you are in a position to know that the outcome that would actually result were you to φ is better than the outcome that would result from ψing. Or, for example, if ψ is predominated over by φ, then you are in a position to know that ψ isn't better than φ and that, in fact, ψ might be worse.

Furthermore, it's plausible to think that what you're in a position to know about the value-relations that actually hold between your options is relevant to what it is, and is not, rationally permissible to do. For example, if you know that the outcome that would actually result from performing φ is better than the outcome that would result from performing ψ, it is rationally impermissible to ψ. Rationality is about doing what makes the most sense given your perspective. And it makes no sense to ψ, when you could φ instead, if you know that φ is better than ψ in the actual world.

In Vacation Boxes, taking the Regular box is no worse than taking the Larger box. And so *Prospectism* violates the following "dominance" principle:[6]

[6] The point that *Prospectism* violates such a principle is made in [Bales et al., 2014], who call the principle *Competitiveness*. Rabinowicz [2016] makes a very similar point, but calls the principle "complementary dominance." And, although not explicitly put in terms of "dominance," Hare [2010],

[The Never Worse Principle]

If, for all available options ψ, φ is no worse than ψ, then it's rationally permissible to φ.

Partition the ways the world might be into two: the worlds in which the coin landed heads (H), and the worlds in which the coin landed tails (T). Because you don't prefer A^+ to B, you don't prefer $(L \wedge H)$ to $(R \wedge H)$; because you don't prefer B^+ to A, you don't prefer $(L \wedge T)$ to $(R \wedge T)$. So, according to *The Never Worse Principle*, it's permissible to take the Regular box. But, as we've seen, *Prospectism* entails that it is *not* permissible to take the Regular box.

But should we accept *The Never Worse Principle*? In the next section, we'll look at two arguments for it. The first argument supports the principle by appealing to an analogy with the more traditional dominance principles. The second argument appeals to what you are in a position to know about the value-relations between your options, and what this knowledge means for rational choice. I don't think either of these arguments are entirely successful (although I am sympathetic to the second of the two) and will offer a response to each on behalf of the *Prospectist*. However, I will go on to show that *Prospectism* violates a principle even weaker, and better supported, than *The Never Worse Principle* in virtue of the fact that it will sometimes require you to perform options that are *predominated* (i.e., that are guaranteed to be no better, and might be worse, than another).

4 *The Never Worse Principle* and Permissibility

Argument 1: the Analogy. Let me present a defense of the principle that I think is, ultimately, unpersuasive. The argument appeals to an analogy. It goes like this.[7] *The Never Worse Principle* is analogous to the following principle:

---

too, makes this point.

[7] This defense is very closely related to the argument that Bales et al. [2014] offer for *Competitiveness*. They appeal to an analogy: *Competitiveness* is the basis of a principle (which they call *Strong Competitiveness*) that is the analogue of the principle of weak dominance, a traditional dominance principle familiar from decision and game theory. Recall: an option φ *weakly dominates* option ψ if, for every way the world might be, φ's outcomes are as good as or better than ψ's outcomes and there's some way the world might be according to which φ's outcome is better than ψ's. If φ weakly dominates all other available options, the you rationally ought to φ. That's *The Principle of Weak*

[The Always As Good Principle]

If, for all available options ψ, φ is at least as good as ψ, then it's rationally permissible to φ.

In fact, in the absence of incomplete preferences — that is, when $X \not\prec Y$ entails $X \succeq Y$, and *vice versa* — the two principles are logically equivalent. Furthermore, *The Always As Good Principle* is surely correct.[8] And, absent some reason for thinking that the two principles are not analogous, we have good reason to accept *The Never Worse Principle* as well.

The problem with this argument is that we have very good reason to think that the two principles are not analogous. The reason, in brief, is that parity is not indifference. In order to see the disanalogy, it will be helpful to look at an argument for *The Always As Good Principle* and, then, show why an analogous argument cannot be made to support *The Never Worse Principle*.

The argument goes like this. Suppose that φ is at least as good as all the other available options. Then, distinguish between two cases:

(1) φ is at least as good as all the other options, and none of the other options are at least as good as φ.

(2) φ is at least as good as all the other options, and there are some other options that are at least as good as φ.

Dominance. Bales et al. [2014] take this to be "[o]ne of the least controversial principles of rational choice [. . .]" (pg. 459) and, if we replace "as good as or better than" in the principle of weak dominance with "not worse than," we arrive at what Bales et al. [2014] call *Strong Competitiveness*, which they consider to be to be at least as plausible as *The Principle of Weak Dominance* (pg. 460). Because *Competitiveness* says "if an option is competitive, then it is rationally *permissible* to take it," Bales et al. [2014] consider it to be a "simpler, and even more compelling, principle" than *Strong Competitiveness* (pg. 460).

There are a number of problems with this argument. First, *Strong Competitiveness*, in addition to being logically stronger than *The Principle of Weak Dominance*, is false (see fn. 18). So the former certainly isn't at least as plausible as the latter. Second, *Strong Competitiveness* and *Competitiveness* are logically independent — even if the former were true, the latter needn't be — so it's not clear why we should find it "even more compelling."

[8] Or at least it is if it's understood to apply only to cases in which your options are suitably independent of the states of the world. Bales et al. [2014] restrict their discussion of dominance principles to cases in which your options and the states are *probabilistically* independent. I think that's overkill: *causal* independence is enough. But, for the time being, I will follow their lead.

Consider case (1). If φ is at least as good as all the other options, and none of the other options are at least as good as φ, then φ weakly dominates all other available options.[9] And it's plausible to think (although I won't argue for it here) that if some option weakly dominates all the others, then you ought to take it. And so it's, at the very least, permissible to.

Consider case (2): φ is at least as good as all the other options, and there are some other options that are at least as good as φ. Let ψ be one such option. If φ is at least as good as ψ and ψ is at least as good as φ, then you're indifferent between their potential outcomes in each state. In other words, the two options pay-out the same *utility* in each state. Because rationality is about doing what's best given your beliefs and preferences, φ and ψ are effectively the same in all the respects that should matter to rationality. Therefore, φ should be rationally permissible if ψ is and *vice versa*; and likewise for all of the options that are at least as good as φ. But, drawing from the argument in case (1), you're required to take one of those options. And so — because they are all effectively the same — it's permissible to take any one of them. And so it's permissible to φ.

The analogous argument in support of *The Never Worse Principle* is not nearly as compelling, though. Suppose that φ is no worse than any of the other available options. Again, we can distinguish between two cases:

(1′) φ is no worse than the other options, and none of the other options are no worse than φ.

(2′) φ is no worse than the other options, and there are some other options that are no worse than φ.

Consider case (1′). If φ is no worse than the other options, and none of the other options are no worse than φ, although it doesn't follow that φ weakly dominates

[9] Here's why. If φ is at least as good as all the other options, then, for all $S \in \mathbf{S}$ and all available options ψ, $(\varphi \wedge S)$ is weakly preferred to $(\psi \wedge S)$. If none of the other options are at least as good as φ, then, for each of the other available options ψ, there will be some state S^* such that $(\psi \wedge S^*) \not\succeq (\varphi \wedge S^*)$. And because φ is at least as good as ψ, $(\varphi \wedge S^*) \succeq (\psi \wedge S^*)$. So, it must be that $(\varphi \wedge S^*) \succ (\psi \wedge S^*)$. Therefore, for all states S and every other available option ψ, you weakly prefer $(\varphi \wedge S)$ to $(\psi \wedge S)$ and there is, for each of those alternatives, some state S^* such that you *strictly* prefer $(\varphi \wedge S^*)$ to $(\psi \wedge S^*)$. In other words, φ weakly dominates all other available options.

all the other options, it is true that φ predominates over them.[10] And one might think (although I won't argue for this here) that if φ predominates over all available alternatives, it's rationally permissible to φ. The problem arises for case (2′). Suppose that φ is no worse than ψ and that ψ is no worse than φ. This might be because, according to your preferences, the two are at least as good as each other. In which case, following (2) above, the two options are effectively the same. And so the requirements of rationality should make no distinction between the two. But it might rather be — as it is in Vacation Boxes — that you regard their outcomes, in each state, to be on a par. And it's far less obvious that rationality should make no distinction between two options when you regard their outcomes, in each state, to be on a par. Parity is not indifference. And, perhaps, when there are multiple options that are no worse than the others, you should take the one(s) with the best prospects?

All this is *not* to say that *The Never Worse Principle* is incorrect. Rather, the point is that it's not *obviously* correct and that the argument by analogy is too quick. Let's look at a better argument.

Argument 2: the role of rationality, and what you know about the value-relations that hold between your options. As briefly alluded to in §3, if φ is *no worse than* ψ, you are thereby in a position to know something about how the actual values of those options compare: namely, that the outcome that would actually result from φing isn't worse than the outcome that would result from ψing. If you think that knowledge about the value-relations actually holding between your options constrain what rationality requires of you, in the way the following principle holds, then *The Never Worse Principle* follows straightaway:

> [Known Value-Relations]
>
> If you know that φ is actually more valuable than all other options, you are rationally required to φ. If you know that none of the other options are actually more valuable than φ, then it's rationally permissible to φ.

[10] Here's why. If φ is no worse than the other options, then, for all states S and all other options ψ, $(\varphi \wedge S) \not\prec (\psi \wedge S)$. And if none of the other options are no worse than φ, then, for each ψ, there is some state S^* such that $(\psi \wedge S^*) \prec (\varphi \wedge S^*)$. And so φ predominates over all other options.

Schoenfield [2014] claims that traditional, single-function expected utility theory satisfies this principle and then goes on to argue that any generalization of expected utility theory, in order to be adequate, must satisfy it as well.[11] The idea is that we must accept something like Known Value-Relations in order for decision theory to, correctly, play the role it was intended to play.[12] Why? Because, ultimately, we should be concerned with bringing about valuable outcomes. And a decision theory that says that you should φ, rather than ψ, when you know that φing won't result in an outcome any more valuable than the one that would result from ψing is offering advice that goes beyond what is of concern: value. If you know that φing won't net you more value than ψing, what reason could there be to take it?

This is a compelling argument, but I want to raise two related worries. First, it's not true that traditional, single-function expected utility theory satisfies Known Value-Relations (or, at least, whether or not it does is controversial). Evidential decision theory, which is a kind of expected utility theory, violates Known Value-Relations in Newcomb cases.[13] In the Newcomb Problem, you are in a position to know that *Two-Boxing* would net you more value than *One-Boxing*. But, because *One-Boxing* provides you with significant evidence that you'll receive a million dollars and *Two-Boxing* provides you with significant evidence that you

[11] Schoenfield [2014] calls this principle Link because it links facts about *expected value* to what is known about *value*. Link is less general than the principle presented here — it's stated in terms of two available options, and it's restricted to "cases in which considerations of value are the only ones that are relevant" — but these are superficial differences that won't matter for our purposes.

[12] Schoenfield [2014] says "if Link is rejected, expected value theory cannot play the role that it was intended to play: namely, providing agents with limited information guidance concerning how to make choices in circumstances in which value-based considerations are all that matter." (pg 268). But, of course, this isn't *literally* true. *Prospectism* rejects Link and, yet, it *does* provide guidance to agents with limited information. The issue is whether the guidance it provides is *correct*. Schoenfield [2014] thinks it's not — the view, in virtue of violating the second clause of Link, "is imposing requirements that transcend what we actually care about: the achievement of *value*" (pg 268) — but, as we'll see, this isn't obvious.

[13] The Newcomb Problem was first discussed in print by Nozick [1969], who attributes it to the physicist William Newcomb. Here's the case. Before you, there are two boxes: an opaque box, which either contains a million dollars or nothing; and a transparent box, which contains a thousand dollars. You have the option, either, to take only the opaque box (*One-Box*) or to take both the opaque and the transparent box (*Two-Box*). Here's the catch. Whether the opaque box contains the million dollars or nothing has been determined by the prediction of a super-reliable predictor. If the predictor predicted that you'd *One-Box*, she put a million dollars in the opaque box; if she predicted that you'd *Two-Box*, she put nothing in the opaque box.

won't, evidential decision theory recommends *One-Boxing* over *Two-Boxing*. Evidential decision theory might be wrong (and, personally, I think it is), but the fact that there is a version of traditional, single-function expected utility theory that violates Known Value-Relations undermines Schoenfield [2014]'s contention that such a principle must be accepted in order for decision theory to, correctly, play the role it was intended to play. At best, the argument shows that Known Value-Relations is central to one *particular* conception of instrumental rationality (the one underlying causal decision theory, for example).[14] Furthermore, given that Known Value-Relations holds only relative to some ways of thinking about instrumental rationality, it's open to the *Prospectist* to claim that they, like the *Evidentialist*, are thinking about instrumental rationality differently.

Is there a plausible conception of instrumental rationality supporting *Prospectism's* verdict in Vacation Boxes? Hare [2010, 2013] articulates one. Here, roughly, is the idea. Being instrumentally rational is about doing what you have the most reason to do. In Vacation Boxes, you have a reason to take the Larger box over the Regular box (you'll get a dollar) and you have *no* reason to take the Regular box over the Larger box (anything that can be said in favor of the former can equally well be said in favor of the latter). Therefore, you have more reason to take the Larger box than the Regular box. And — because instrumental rationality requires you to do what you have the most reason to do — you are rationally required to take the Larger box.[15]

Let me pause, briefly, to illustrate the argument with an example, which I think will help throw the conception of rationality underlying it into sharper relief. Suppose you know that, after making your decision in Vacation Boxes, you'll be asked

[14] In fact, if we assume your preferences can be represented with a utility-function, causal decision theory entails Known Value-Relations. And, if we define the *actual value* of an option φ, $V_@(\varphi)$, to be the utility you assign to the outcome that would result were you to φ, then the causal expected utility of φ equals your best estimate of φ's actual value: that is, $U(\varphi) = \sum_v Cr\left(V_@(\varphi) = v\right) \cdot v$.

[15] This is one of the arguments Hare offers in favor of *Prospectism's* verdict in cases like Vacation Boxes (see [Hare, 2010, 2013]). The argument is criticized by Bales et al. [2014] and Schoenfield [2014], largely on the grounds that it fails to appreciate the potentially complicated ways in which reasons can interact (e.g., "Reasons interact in complex ways and they don't always add up as one might expect them to." [Schoenfield, 2014, pg. 273]). All parties want to accept that rationality requires you to do what you have the most reason to do, but I think, if you want to resist Hare's argument, this is untenable. Presenting the argument for why, though, would take us too far afield.

to explain why you decided as you did. If you choose to take the Larger box, there is something you can say: namely, "I knew that I would get a dollar if I took the Larger box, but that I wouldn't get a dollar if I took the Regular box." If you choose the Regular box, however, there is nothing satisfying that you can say to explain your choice. You can't say, for example, that you chose as you did because you might get prize *A* (or because you might get prize *B*) because this isn't a consideration that speaks in favor of taking the Regular box *rather than* the Larger box, given that you know the Larger box might contain prize *A* (and might contain prize *B*) as well. Likewise, you can't say "I knew there was a 50% chance of getting *A*," or "I knew I would get a prize that's uniquely good," or "I knew the prize I would get wouldn't be worse than the one I would get had I chosen otherwise." None of these considerations distinguish between the two boxes, and so none of them can be used to explain why you chose the one over the other.

There are, of course, features that distinguish between the two, but these features, given what we're assuming you care about, won't make for a *satisfying* explanation. For example, you can't say "I knew that I would get a prize that was inside a regular-sized box if I took the Regular box, but that I wouldn't if I took the Larger box," because you care about the prizes, not the sizes of the boxes they come in.

Furthermore, while *there are* reasons for taking the Regular box over the Larger box, none are reasons that you *have* (at least prior to discovering which prizes are in which box). For example, if, as a matter of fact but unbeknownst to you, the Regular box contains prize *B*, then all the uniquely valuable things about *B* are reasons to take the Regular box over the Larger one. But, because you don't know which box contains *B*, you aren't in a position to cite these reasons in an explanation of your choice.

Given that there's nothing you can say to explain choosing the Regular box over the Larger but that there is something you can say to explain choosing the Larger box, if you want to be able to explain your behavior, you shouldn't choose the Regular box. But what is rationality about if not behaving in ways the make sense — that are explicable — in light of your perspective?

This conception of rationality allows for violations of KNOWN VALUE-RELATIONS because it holds that rationality is about doing what you have the most reason to do, and, in cases like Vacation Boxes, you have a reason to take the Larger box over the Regular box but no reason to take the Regular box over the Larger box. The fact that the Regular box is no worse than the Larger box isn't enough to render the taking of it rationally permissible because, in this case, that fact doesn't constitute a consideration that speaks in favor of taking it rather than taking the Larger box. And, so the thought might go, it's permissible to take an option only if something can be said in favor of doing so (unless nothing can be said in favor of doing anything; in which case, everything is permitted).

It'll often be the case that knowing that φ isn't worse than the other options *will* provide you with a reason to φ, but not always. And so — at least, offhand — the *Prospectist* can accept that knowledge about the value-relations that actually hold between your options is relevant to what rationality requires of you — and thus accept much of what Schoenfield [2014] says about the connection between rationality and value — and yet deny that such knowledge is relevant in the way Known Value-Relations claims it to be.[16] Knowing what value-relations actually hold between your options is relevant to what rationality requires of you only insofar as that knowledge affects what reasons you have. And so, while this conception of rationality violates Known Value-Relations, it can endorse the connection between rationality and value that's captured by the following, weaker, constraint:

> [Weak Link]
>
> If you know that none of the other options are actually more valuable than φ, and you have some reason to φ, then it's rationally permissible to φ.

[16] Hare [2013, pg. 51] defends *Prospectism* against arguments like Schoenfield [2014]'s along these lines. He gives an explanation for why these arguments might seem attractive even though they are, in his opinion, unsound: sometimes, when we learn that φ is no worse than ψ, it ceases to be true that we have a reason to ψ and no reason to φ. Much depends on *exactly what* we know about the value-relation that holds between our options. For example, if you know that φ is no worse than ψ because you know that φ is actually better than ψ, then you'll have a reason to φ rather than ψ (and, presumably, a decisive reason at that). If you know that φ is no worse than ψ because you know that φ and ψ are actually equally good, then you won't have more reason to ψ than to φ (either because you won't have any reason to do one rather than the other, or because you will have a reason to φ that perfectly balances your reasons to ψ). But if you know that φ is no worse than ψ because you know that φ and ψ are actually on a par, you might, like in Vacation Boxes, have a reason to ψ and no reason to φ. As I'll argue in the next section, however, this defense isn't entirely adequate.

On this view, the reason it's impermissible to take the Regular box in Vacation Boxes, in spite of the fact that you know it isn't actually worse than taking the other box, is that you don't have a reason to do so. But were you to have one — if there were something you could say in favor of taking it rather than the other — shouldn't it be permissible to take it (at least, if you also know that doing so won't be worse than taking the other)?

Let's grant that rationality is about doing what you have the most reason to do. Even so, Schoenfield [2014] is surely right: the requirements of rationality must make sense given what you know about the value-relations that actually hold between your options. The thought is that, if you know that φ is no worse than ψ, the only thing that can prevent it from being permissible to φ is that you lack a reason to do so. But if you have a reason to φ, and you know that φing won't result in a worse outcome, it's permissible to do so.

But Weak Link cannot be used, as Known Value-Relations was, to support *The Never Worse Principle*. If φ is no worse than the other available options, you are in a position to know that none of the other options are actually more valuable than φ. But that's compatible with you having no reason to φ. And so Weak Link, unlike Known Value-Relations, doesn't entail *The Never Worse Principle*.

Taking Stock and Looking Ahead. We've looked at two arguments for *The Never Worse Principle*. I've argued that neither is entirely successful. The first argument appealed to an analogy between *The Never Worse Principle* and *The Always As Good Principle*. But, as we saw, there is an important disanalogy — which turns on the difference between indifference and parity — between the respective arguments that can be made in support of these two principles. The second argument appealed to a principle, Known Value-Relations, which connects the requirements of rationality to what you are in a position to know about the value-relations that actually hold between your options. I argued that Known Value-Relations is true only if we think about instrumental rationality in a particular way. And that this way isn't how *Prospectists* are thinking about rationality, nor need it be. Furthermore, the *Prospectist* can seemingly accept, as the argument for Known Value-Relations contends, that there is a tight connection between the requirements of rationality and what's known about the value-relations holding between your options by accepting Weak Link (and emphasizing the ways in which such knowledge, at least typically, affects the reasons you have).

We don't have a satisfactory argument for *The Never Worse Principle*. But *Prospectism* is not yet in the clear. In the next section, I will show that the view violates a logically weaker dominance principle — what I will call *The Principle of Predominance* — to which the worries raised above do not apply. As we'll see, *Prospectism* will sometimes forbid you from taking an option even though you know it's no worse, and think it's likely to be better, than your other options. This shows that *Prospectism* cannot be supported by any conception of rationality that endorses Weak Link. And therefore it divorces what you should do from what you know about the relevant value-relations in a way that might seem implausibly radical.

5 *Predominance*

Consider a variant of Vacation Boxes: **Pay or Roll**. Everything is the same as before, except this time your options are slightly different: you can pay a small fee for the Larger box (L^-), or you can roll a fair six-sided die to decide which of the two boxes you get (M, for "mixed option"). If the die rolls a 1, you get whichever prize happens to be in the Regular box; if it rolls anything other than a 1, you get whichever prize happens to be in the Larger box.

- **L^-** You pay $\$\varepsilon$ for the Larger box.
- **M** A six-sided die is rolled. If it rolls a 1, you get the Regular box; if it rolls 2–6, you get the Larger box.

Taking the Larger box has better prospects than option M. And if we pick a suitably small enough value for $\$\varepsilon$, L^- will too. Making a few simplifying assumptions about your preferences — namely: that you value small sums of money linearly; that you value these small sums of money independently of the vacation prizes; you are, if risk-seeking, only slightly so; and that you care about the prizes and money, not the ways in which you might receive them — any amount less than 15¢ will do.[17] Because paying 15¢ for the Larger box has better prospects than taking

[17] If you do not value money linearly — if, for example, receiving 85¢ is less than five-sixths as good than receiving \$1 — we'll have to pick a smaller value for $\$\varepsilon$. In particular, we should pick an

option M for free, *Prospectism* says that you're rationally required to take L^- over M; it is rationally impermissible, according to *Prospectism*, to choose M.

Pay or Roll

Coin Toss:	Heads		Tails	
Die Roll:	1	2–6	1	2–6
Option L^-	$A^{+85¢}$	$A^{+85¢}$	$B^{+85¢}$	$B^{+85¢}$
Option M	B	$A^{+\$1}$	A	$B^{+\$1}$

But is it irrational to choose M over L^-? Here's a reason to think it is not. Much like the Regular box in Vacation Boxes, M is guaranteed to be no worse than the other available option. But, in addition, there are some ways the world might be in which it does better. In other words, M predominates over L^-. And you might think: if an option predominates over all others, it's rationally permissible to take it.[18]

amount so that $u(\$1 - \varepsilon) > \frac{5}{6} \cdot u(\$1)$. If you're fairly risk-seeking — see [Buchak, 2013] for a way of modeling agents who are genuinely risk-averse and genuinely risk-seeking — then, again, we'll have to pick a smaller value for $\$\varepsilon$. In particular, making use of Buchak [2013]'s *Risk-Weighted Expected Utility Theory*, where $0 \leq r_x(p) = p^x \leq 1$ is a risk-function representing your attitude to risk, L^- will have better prospects than M if $r\left(\frac{5}{12}\right) + r\left(\frac{11}{12}\right) - r\left(\frac{1}{2}\right) \leq \frac{u(\$1-\varepsilon)}{u(\$1)}$. In order for M to come out ahead of L^-, when $\varepsilon \leq 15¢$, you'd need to be more than just slightly risk-seeking. If you're enough of a risk-seeker, might there no value small enough for *Prospectism* to recommend choosing L^- over M? Yes. You could be such a risk-seeker that you're disposed to avoid sure-things at all costs. This is an extreme — and arguably irrational — way to be. It should provide only cold comfort for the *Prospectist* then.

[18] It might be tempting to think that, if an option predominates over all others, then you're rationally *required* to take it. (This is entailed by — and, when there are only two options at play, equivalent to — what Bales et al. [2014] call *Strong Competitiveness*). But resist the temptation because *that* principle is surely false. For example, it would require you to φ over ψ — even if you think it's overwhelmingly likely that the two are on a par — so long as there is some chance, no matter how small, that φ might be ever-so-slightly better than ψ. That strikes me as implausible. But even more seriously: imagine a case in which there are more than two options such that each of which is predominated by one of the others. In such a case, this principle will say, when considering the options pair by pair, that you are rationally required to choose in a way that will result in cyclic — and, hence, *strongly* money-pumpable — choice behavior. I think this is a decisive reason to reject *Strong Competitiveness*. (There is, it might be objected, a similar money-pump worry for the weaker principle endorsed in the main text: the same cases will be ones in which, according to that principle, it's not impermissible to be money-pumped. This is a much less serious problem, though, because, rather than *forcing* you into a sub-optimal outcome, the principle merely fails at

[The Principle of Predominance]

If, for all other options ψ, φ predominates over ψ, then it's rationally permissible to φ.

Prospectism, in virtue of recommending L^- over M, violates this principle. If *Prospectism* is correct, then there are cases — this one, for example — in which you are *rationally required* to choose an option that can do no better, and is likely to be worse, than some other.

The violation of *The Principle of Predominance* presents a more serious challenge to *Prospectism* than its violation of *The Never Worse Principle*, and not only in virtue of the former being logically weaker than the latter. The fact that *Prospectism* violates *The Principle of Predominance* means that the view cannot be supported by the conception of rationality — or any conception of rationality satisfying Weak Link — that was sketched, in the previous section, as a response to Schoenfield [2014].

Here's why. Because M predominates over L^-, you are in a position to know that the outcome that would result from taking M is not worse than the one that would result from taking L^- and, furthermore, you think it's likely to be better. Therefore, unlike the Regular box, there *are* things to be said in favor of taking option M. For example, you know that if you take M you're likely to get a dollar and that you definitely won't get a dollar if you take L^- (you'll get only 85¢ instead).[19]

preventing you from stumbling into one. In the former (more serious) case, there's nothing you can do to satisfy what's rationally required of you; in the latter (less serious) case, there is: it's rationally permissible for you to turn down any or all of the trades. Moreover, so long as your preferences are incomplete, you are *already* vulnerable to money-pumps of this sort: e.g., it's permissible to trade A^+ for B, it's permissible to trade B for A, but A^+ is better than A. So, these weaker money-pumps aren't reason enough to reject *The Principe of Predominance*, whereas the stronger ones are reason to reject *Strong Competitiveness*. Thanks to an anonymous referee for helpful discussion on these points.)

[19] One might object that, because L^- comes with a sure-thing 85¢ and (per the assumptions we made earlier about your preferences) you'd rather have a sure-thing 85¢ than a five-sixths chance at getting $1, this isn't really a consideration that speaks in favor of taking M over L^- after all. To bring out the thought behind this objection, consider the following example. Suppose that if you φ you'll get a dollar and that if you ψ you'll get two. Is the fact that you'll get a dollar if you φ and that you won't get *a* dollar if you ψ a reason to take φ over ψ? It's unclear. You might think: *no* because you might think getting *two* dollars entails getting *one* and so "I'll get a dollar" doesn't properly distinguish between the two options. But what about "I'll get *exactly* one dollar"? That distinguishes between

Or — what strikes me as an even more compelling consideration in favor of taking M — that doing so will likely result in an outcome that you prefer. The same cannot be said for taking L^-: you are certain that taking it will *not* result in an outcome you prefer. And so you clearly have reason to choose M. In Vacation Boxes, it was the fact that you had *no reason* to take the Regular box that was meant to justify the impermissibility of doing so (given that you knew its outcome would be no worse than the other). But in this case, you know M is no worse *and* you have reason to take it!

The conception of rationality (sketched in the previous section), which holds that rationality is about doing what you have the most reason to do, but which also countenances a tight link — albeit, one mediated by one's reasons — between the requirements of rationality and what's known about the options' value-relations, doesn't support *Prospectism* after all. The *Prospectist* can retain the first part — that rationality is about doing what you have the most reason to do — but it must give up the second part. The *Prospectist* can, and should, say:

> "Although you *do* have reason to take M over L^-, you also have reason to take L^- (you're sure to get 85¢!). And, on balance, your reasons favor L^- over M. And so — because rationality is about doing what you have the *most* reason to do — you are required to choose L^-."

But why think the balance of reasons weigh in favor of L^- over M? You have a reason to take L^- over M: it's that you're sure to get 85¢ if you do. And, because you prefer a sure-thing 85¢ to a five-sixths chance at a dollar, that reason might very well outweigh the fact that if you take M you're likely to get a dollar and definitely won't otherwise. But, as mentioned before, that's only *one* of the considerations that speaks in favor of choosing M. The other considerations concern what you know about the value-relations that actually hold between your options:

---

the two. But you might think *that's* not a reason either; it doesn't speak *in favor* of taking φ given that getting exactly one dollar isn't a good thing when compared to getting two. And, similarly, a five-sixths chance at a dollar is not a good thing when compared to a sure-thing 85¢. Alternatively, you might think: no, "I'll get (exactly) a dollar" *is* a reason to φ rather than ψ; it's just that, in this case, that reason is clearly *outweighed* by the fact that you'll get two dollars if you take ψ. I think that's the right thing to say, and that's what we should say about Dice Roll: the fact that you're likely to get a dollar if you take M (and that you won't get a dollar if you take L^-) is a reason to choose M over L^-, even if it's a reason that's ultimately outweighed by others.

r_1 If I take M, it's likely that I'll get a prize that's better than the one I would've gotten had I chosen otherwise, but I'm not likely to get a better prize if I take L^- (in fact, I know that if I take L^-, I definitely *will not* got a prize that's better than the one I would've gotten by choosing differently).

r_2 If I take M, I won't get a prize that's worse than what I would've gotten by choosing differently, but I might get a worse prize if I take L^-.

The *Prospectist* must maintain, either, that r_1 and r_2 aren't genuine reasons to choose M or that, if they are, they are outweighed by the promise of that 85¢ you're sure to get by choosing L^- instead. I ultimately don't think either option is plausible. First, it's hard to see why r_1 and r_2 wouldn't count as genuine reasons: both sound like plausible things you could say to justify choosing M over L^-. The better option, then, is to say that the reasons are outweighed. But because *Prospectism* evaluates options solely in terms of those options' corresponding prospects (and because an option's corresponding prospects abstract away from which of its outcomes reside in which states), reasons — like r_1 and r_2 — that concern the value-relations holding between your options will *always* be outweighed. And so reasons like these appear to be weight*less*. But a weightless reason is no reason at all!

Here's an example to bring this out. Suppose we sweeten the Regular box with a gamble that pays out $1 with a five-sixths chance. Call that option R^*. Notice that R^* has the same prospects as M. Let ε be the amount such that you are indifferent between a sure-thing $\$(1 - \varepsilon)$ and a five-sixths chance at getting $1. Suppose we sweeten the Larger box, not with a dollar, but with $\$(1 - \varepsilon)$ instead. Call this option L^*. Option M predominates over L^* while R^* does not.

Coin Toss:	Heads		Tails	
Die Roll:	1	2–6	1	2–6
Option L^*	$A^{+\$(1-\varepsilon)}$	$A^{+\$(1-\varepsilon)}$	$B^{+\$(1-\varepsilon)}$	$B^{+\$(1-\varepsilon)}$
Option R^*	B	$B^{+\$1}$	A	$A^{+\$1}$
Option M	B	$A^{+\$1}$	A	$B^{+\$1}$

According to *Prospectism*, it's rationally permissible to take either L^* or R^*.[20] So the reason(s) you have for taking L^* over R^* must perfectly balance the reason(s) you have for taking R^* over L^*. Similarly, according to *Prospectism*, it's rationally permissible to take either L^* or M. And so, also, the reason(s) you have for taking L^* over M must perfectly balance the reason(s) you have for taking M over L^*. Option M has the same prospects as option R^*, so the reasons that favor R^* over L^* also favor M over L^*. But, unlike the choice between L^* and R^*, when choosing between L^* and M, you know something about the value-relations holding between the options which provides you with additional reasons to favor M over L^*. But, as this example brings out, for *Prospectism*, these sort of reasons can have no effect on an option's permissibility. They are inert.

But that's implausible. How can something be a reason — and one that, as the above example makes clear, is distinct from your other reasons — and yet carry no weight in any circumstance? Reasons that cannot, even in principle, make an otherwise impermissible option permissible or an otherwise non-required option required don't seem like reasons at all! Furthermore, following Schoenfield [2014], you might think that these kinds of reasons — ones concerning the value-relations holding between your options — are the ones that matter most. Ultimately, it's how the actual values of your options compare that you care about, not your options' prospects. And yet *Prospectism* effectively ignores these reasons entirely.

6 Conclusion

Prospectism conflicts with dominance reasoning. It sometimes requires you to do something that's no better than the alternatives and might (or even likely!) be worse. In light of this, if *Prospectism* adopts a conception of rationality according to which you're required to do what you have the most reason to do, then it must also adopt a view about reasons that is, at best, controversial and, at worst, implausible. In particular, *Prospectism* gives no weight to reasons concerning the value-relations that actually hold between your options. But that's to treat these reasons

[20] Here, like before, I am making some assumptions about your preferences. If those assumptions don't hold, we can change the example so that they do.

as if they aren't reasons at all. I think this is a significant mark against views, like *Prospectism*, that evaluate options solely in terms of their corresponding prospects.

But if we are to reject *Prospectism*, what are the alternatives? This is a good question. And, for reasons that are outside the scope of this paper, I'm not particularly optimistic that there are *any* general decision theories involving parity that are lacking in significant drawbacks.[21] It might come down to which kinds of drawbacks one can tolerate. (And perhaps these alternative views are themselves on a par.) But whatever the case may be, *Prospectism* has some significant problems, which are, in my opinion, difficult to tolerate.

References

A. Bales, D. Cohen, and Toby Handfield. Decision theory of agents with incomplete preferences. *Australasian Journal of Philosophy*, 92(3):453–470, 2014.

Lara Buchak. *Risk and Rationality*. Oxford University Press, 2013.

Ruth Chang. The possibility of parity. *Ethics*, 112(4):659–688, 2002.

J. Dubra, F. Maccheroni, and E. Ok. Expected utility theory without the completeness axiom. *Journal of Economic Theory*, 115(1):118–133, 2004.

[21] Here's one example. Consider a view that endorses Weak Link: it agrees with *Prospectism* in cases like Vacation Boxes (you're required to take L over R) but disagrees with *Prospectism* in cases like Pay or Roll (it's permissible to take either option). A view like this offers counterintuitive recommendations concerning cases of *probabilistic sweetening*. Imagine a variant of Vacation Boxes in which L has been sweetened with \$2. You have the opportunity, before choosing, to sweeten R either with a dollar ($R^{\$1}$) or with a lottery ticket that pays out a million dollars on the very slim chance that it wins (R^{ℓ}). The chance of the ticket winning is so low that you prefer the dollar, all else equal. Because $R^{\$1}$ never does better than L and you have no reason to take it, you should prefer L to $R^{\$1}$. However, because R^{ℓ} never does worse than L and might (with very small probability) be better, you shouldn't prefer L to R^{ℓ}. This is, at the very least, odd: sweetening R with something you prefer appears to make it worse! Moreover, you have intransitive preferences: you should prefer L to $R^{\$1}$, you should prefer $R^{\$1}$ to R^{ℓ}, but you shouldn't prefer L to R^{ℓ}. Offhand, this seems like a good reason to reject this view in favor of one that disagrees with *Prospectism*'s recommendation in Vacation Boxes as well. However, *those* views have problems of their own. A full discussion of the strengths and weaknesses of the alternative views, though, are outside the scope of this paper.

Ozgur Evren and Efe Ok. On the multi-utility representation of preference relations. *Journal of Mathematical Economics*, 47:554–563, 2011.

Tsogbadral Galaabaatar and Edi Karni. Subjective expected utility theory with incomplete preferences. *Econometrica*, 81(1):255–284, 2013.

I.J. Good. Rational decisions. *Journal of the Royal Statistical Society*, B(14):107–114, 1952.

Caspar Hare. Take the sugar. *Analysis*, 70(2):237–247, 2010.

Caspar Hare. *The Limits of Kindness*. Oxford University Press, 2013.

James Joyce. A defense of imprecise credences in inference and decision making. *Philosophical Perspectives*, 24(1):281–323, 2010.

Isaac Levi. Imprecision and indeterminacy in probability judgment. *Philosophy of Science*, 36:331–340, 1985.

Isaac Levi. *Hard Choices*. Cambridge University Press, Cambridge, 1986.

Isaac Levi. Why rational agents should not be liberal maximizers. *Canadian Journal of Philosophy*, 38(Supplementary Vol 34):1–17, 2008.

Robert Nozick. Newcomb's problem and two principles of choice. In N. Rescher et al., editor, *Essays in Honor of Carl G. Hempel*. Reidel, 1969.

E. Ok, P. Ortoleva, and Gil Riella. Incomplete preferences under uncertainty: Indecisiveness in beliefs versus tastes. *Econometrica*, 80(4):1791–1808, 2012.

Wlodek Rabinowicz. Incommensurability meets risk. unpublished, July 2016.

Susanna Rinard. A decision theory for imprecise credences. *Philosopher's Imprint*, 15(7):1–16, 2015.

Miriam Schoenfield. Decision making in the face of parity. *Philosophical Perspectives*, 28(1):263–277, 2014.

Amartya Sen. Incompleteness and reasoned choice. *Synthese*, 140(1/2):43–59, May 2004.

Brian Weatherson. Decision making with imprecise probabilities. manuscript, November 2008.

Paul Weirich. *Realistic Decision Theory: Rules for Nonideal Agents in Nonideal Circumstances*. Oxford University Press, 2004.