

Hard Choices Made Harder

RYAN DOODY

Forthcoming in *Value Incommensurability: Ethics, Risk, and Decision-Making*, ed. H. Andersson and A. Herlitz.

Abstract

How should you evaluate your choices when you're unsure what their outcomes will be? One popular answer is to rank your options in terms of their expected utilities. But what should you do when you think that the value of their respective outcomes might be incommensurable? In the face of incommensurable values, it no longer makes sense to speak of ranking your options according to expected utility. Are there any general principles to guide us when facing decisions of this kind? If only! This chapter develops an impossibility result: it holds that there are a handful of independently plausible constraints that no such decision theory can jointly satisfy. The result, while depressing, can be used to helpfully classify extant approaches based on which of the constraints they violate.

1 Introduction

Suppose you are deciding between law school and pursuing a PhD in philosophy. This is a hard choice. The life of a philosophy professor—which is what you hope would result from pursuing the degree in philosophy—and the life of a successful lawyer—which is what you hope would result from going to law school—manifest significantly different qualities of what you value in a career. Becoming a successful lawyer promises greater financial security, holidays in Majorca, and a revered

social status. Becoming a philosopher, however, promises greater intellectual satisfaction, lazy mornings, and a great deal of professional autonomy. When it comes to these two careers, you're *ambivalent*: you don't prefer one to the other, nor are you indifferent between the two.¹ It's hard to choose when you're ambivalent.

But this choice is even harder still. For this is also a choice under uncertainty. Merely choosing to *pursue* a particular career doesn't guarantee success. Going to law school *might* result in you becoming a successful lawyer. But it might not. Instead, it might result in you flunking out and buried beneath a mountain of student loans. Similarly, pursuing a PhD in philosophy *might* result in becoming a professor of philosophy, but this is far from guaranteed. Instead, it might result in a very interesting dissertation on the metaphysics of absences and very troubling absences of tenure-track positions in metaphysics. Both options are a *gamble*: how things shake out turn on features of the world you're uncertain about.

When you are uncertain about what will result from your choices, the standard advice is to maximize expected utility (i.e., for each of your options, you should weight the value of its possible outcomes by your credence in that outcome being the one that would actually result were you to perform it. The weighted average of these values is an option's *expected* value. The option with the highest expected value is the one you ought to perform.) But if you're ambivalent between the potential outcomes of your options—as you are in this example—this advice is not helpful. Because you are ambivalent between being a lawyer and being a philosopher, your preferences cannot be represented with a utility-function; and, thus, there is no well-defined quantity whose expectation you can maximize.

What can be said instead? Are there any general principles—or, better still, a full-fledged decision theory—to guide us when facing decisions of this kind? If ambivalence and uncertainty are ubiquitous, which I contend it's plausible to think they are, this is a pressing question.²

¹ In decision theory, this is known as having *incomplete preferences*. And, while the Completeness Axiom, which says that one preferences should be complete, is a mainstay of traditional axiomatic decision theory, it is often, since the very beginning, assumed for the sake of mathematical convenience rather than justified as a genuine rational requirement. See, for example, [von Neumann and Morgenstern, 1944, p. 630], [Savage, 1954, p. 21], [Chernoff, 1954], and [Aumann, 1962].

² Chang [1997, 2002, 2009, 2017] is a notable example of someone who, like me, thinks the phenomenon is ubiquitous. However, I intend 'ambivalence' to refer to a subjective, psychological state. The psychological state is conative, not cognitive: ambivalence might correspond to, but is not identical with, a belief about the objective value-relations that hold between the items. Chang's 'parity', on the other hand, is a (non-transitive) fourth value-relation; in addition to "better than,"

To the end of addressing this question, I develop an impossibility result: there are a handful of independently plausible constraints that no such decision theory can jointly satisfy.

2 The Puzzle of Opaque Sweetening

Let's begin by looking at Caspar Hare's puzzle of opaque sweetening [Hare, 2010, 2013], which itself underlies an impossibility result.

Opaque Sweetening. 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 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 (see Figure 1).

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

Figure 1: Two vacations in opaque boxes, distributed by coin-flip.

Now imagine that \$20 is added to the Larger box. If you choose the Larger box, you will win whichever prize it contains plus \$20. 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.

“worse than,” and “as good as,” two comparable alternatives might be “on a par.” Discussions of related phenomena can be found in [Griffin, 1986, p. 80-98], [Parfit, 1984, p. 431], and [Broome, 2007].

	HEADS	TAILS
L^+	$A^{+\$20}$	$B^{+\$20}$
R	B	A

The puzzle exploits a characteristic feature of being ambivalent—namely, that it, unlike indifference, is *insensitive to mild sweetening*: if you are ambivalent between X and Y , you should likewise be ambivalent between X and Y^+ , a slightly improved version of Y . Let's suppose that you are ambivalent between the adventurous skiing trip in the Alps and the relaxing trip to the beach. There are various good things to be said in favor of each. You don't prefer either one to the other. Moreover, you don't prefer the beach vacation *plus twenty bucks* ($B^{+\$20}$) to the ski vacation, and you don't prefer the ski vacation *plus twenty bucks* ($A^{+\$20}$) to the beach vacation. Small improvements aren't enough to resolve your ambivalence.

Is there a box you're rationally required to take? Or is it rationally permissible to take either? The puzzle brings out a tension between two attractive ideas. One of these ideas supports the verdict that one is rationally required to take, or prefer, L^+ over R . And the other supports the verdict that it is rationally permissible for one to take either.

The first idea is that you ought to evaluate an option in terms of its corresponding *prospects*: the probability-distribution over its potential outcomes. In this case, L^+ 's corresponding prospects (a 50% chance of $A^{+\$20}$ and a 50% chance of $B^{+\$20}$) are better than R 's (a 50% chance of A and a 50% chance of B); and so you ought to prefer, and thus take, L^+ over R .³

The other idea is that rationality shouldn't require you to prefer one option to another if there is no way for the one to be better than the other. In this case, even though you don't know which box contains which prize, you know that you don't prefer what's in L^+ to what's in R . If the coin landed heads, then L^+ contains $A^{+\$20}$

³ Although not always argued for explicitly, there are a number of views committed to this underlying idea: e.g., 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]. The view is also popular among economists: e.g., [Dubra et al., 2004; Evren and Ok, 2011; Galaabaatar and Karni, 2013; Ok et al., 2012]. I'll treat Hare's *Prospectism* as a representative of this family of views [Hare, 2010, 2013]. These views, roughly, represent one's preferences with the *set* of utility-functions that correspond to all of the coherent extensions of one's preference-ordering such that you count as preferring X to Y just in case every function in the set ranks X ahead of Y , you count as being indifferent between X and Y just in case every function in the set ranks them that way, and you lack a preference between the two when there's disagreement among the functions in the set.

and R contains B , and you don't prefer $A^{+\$20}$ to B ; if the coin landed tails, then L^+ contains $B^{+\$20}$ and R contains A , and you don't prefer $B^{+\$20}$ to A . No matter how the coin landed, you're ambivalent between the prizes in each box. There is no way for L^+ to result in an outcome you prefer to what you would get by choosing R instead. And—so the thought goes—you aren't rationally required to prefer, or take, L^+ over R .

This conflict can be presented as an impossibility result: no decision theory, which allows for ambivalence, can satisfy the following four constraints.

NEUTRALITY

If φ and ψ are probability-preserving permutations of each other, you are rationally required to be indifferent between the two.

Two options are *probability-preserving permutations* of each other just in case every outcome that might result from the one is just as likely to result from the other. Here's an example.

	HEADS	TAILS
L	A	B
R	B	A

Given that a fair coin is equally likely to land heads as it is to land tails, L and R are probability-preserving permutations of each other. Every outcome that might result from L —namely, A or B —is equally likely to result from R . According to *Neutrality*, then, you are rationally required to be indifferent between L and R .⁴

The motivation for *Neutrality* is that, when evaluating an option, all that should matter are how good its outcomes are and how likely those outcomes are to result. The identity of the state in which the outcome occurs is only relevant insofar as it

⁴ It might be objected that *really* the two options have different outcomes because, e.g., in the case of L , there is A-when-the-coin-has-landed-heads whereas, for R , there is the outcome A-when-the-coin-has-landed-tails. However, the typical convention in decision theory, which I will be adhering to here, is that outcomes are individuated finely enough to encompass all of the features that matter to the chooser. If the flip of the coin mattered in this way, it would be a mistake to represent, e.g., the outcomes $(L \wedge \text{HEADS})$ and $(R \wedge \text{TAILS})$ with the same symbol. The fact that we *are* representing these outcomes with the same symbols should be taken to apply that HEADS and TAILS are (to borrow a phrase from Ramsey) *ethically neutral*: you are indifferent between $(L \wedge \text{HEADS})$, A , and $(R \wedge \text{TAILS})$, and you are indifferent between $(L \wedge \text{TAILS})$, B , and $(R \wedge \text{HEADS})$.

affects the outcome's value or probability of obtaining. Outcomes are to be individuated in such a way that the states in which they occur don't affect their value. So, options that are probability-preserving permutations of each other are the same in every respect that's relevant to rational choice.

The next constraint is a familiar one. It relies on the notion of *strict dominance*. Let's say that an option φ *strictly dominates* an option ψ just in case, for every state S that might obtain, you prefer φ 's to ψ 's outcome in S .⁵ Informally, we can say that one option strictly dominates another just in case it always does better. (Although we don't need it now, the notion of *weak dominance* will play a role later on. One option *weakly dominates* another just in case, for each state, the former's outcome is at least as good as the latter's, and there is some state in which the former's outcome is better.⁶ Informally, we can say that one option weakly dominates another just in case it's always at least as good and sometimes better.)

STRICT DOMINANCE PRINCIPLE

If φ strictly dominates ψ , you're rationally required to prefer φ to ψ .

The motivation for the *Strict Dominance Principle* is that, if φ strictly dominates ψ , φ is guaranteed to be better than ψ . Although you might not know which outcome would result from performing φ , you do know that, whichever it is, you prefer it to the outcome that would result from performing ψ instead. You know, then, that φ will net you more value than ψ would. Therefore, you should prefer it. Here's an example.

	HEADS	TAILS
L^+	$A^{+\$20}$	$B^{+\$20}$
L	A	B

L^+ strictly dominates L . You don't know how the coin has landed, but you do know that, either way, L^+ will net you a better prize than L would. Because L^+ strictly dominates L , you are in a position to know something significant about how the options compare in the actual world: L^+ is better. Ideally, you'd always prefer the option that promises to be *actually* better. Of course, we don't always know which

⁵ In symbols: φ *strictly dominates* ψ just in case, for all states S , $(\varphi \wedge S) \succ (\psi \wedge S)$.

⁶ In symbols: φ *weakly dominates* ψ just in case, for all states S , $(\varphi \wedge S) \succeq (\psi \wedge S)$ and there is some state S^* such that $(\varphi \wedge S^*) \succ (\psi \wedge S^*)$.

of our options promises to actually be best—but, when we do know (as we do when one of the options strictly dominates the others), it would be irrational not to prefer it.⁷

The third constraint is also a familiar one: *transitivity*. Because we're concerned with a number of different ordering-relations—e.g., strict preference, indifference, ambivalence—there are a number of different corresponding transitivity claims. The first concerns only strict preference:

TRANSITIVITY OF PREFERENCE

If you prefer X to Y and you prefer Y to Z , then you should prefer X to Z .

The next concerns only indifference:

TRANSITIVITY OF INDIFFERENCE

If you're indifferent between X and Y and you're indifferent between Y and Z , then you should be indifferent between X and Z .

And then there are versions of transitivity that mix together preference and indifference, the following of which most concerns us here:

SENSITIVITY TO SWEETENING

If you prefer X to Y and are indifferent between Y and Z , then you should prefer X to Z .

These different claims admit of different possible motivations. The most famous motivations for *Transitivity of Preference*, for example, is the money pump argument: if you violate it, it's possible to construct a series of trades, each of which (if you act on your preferences) it is rational for you to take, but which collectively guarantee a sure-thing loss. Alternatively, you might take transitivity to be partially *constitutive* of the meaning of these notions [Davidson, 1976]. For example, part of what it *means* to be indifferent between Y and Z is that, if you prefer X to

⁷ A quick note of caution. The *Strict Dominance Principle*, like all statewise dominance principles, risks delivering incorrect verdicts when the probability of a state depends on which option is performed [Nozick, 1969]. It's controversial what notion of dependence matters here—a causal or evidential one. Let's hereby avoid the issue by understanding the dominance principles to be restricted to cases in which the states are both causally and evidentially independent of one's options.

Y, you should then prefer X to Z. This is how I propose we interpret *Sensitivity to Sweetening*.⁸

The final constraint has the flavor of a dominance principle. It relies on the notion of one option being *never worse* than another. Let's say that an option φ is *never worse than* an option ψ just in case, for every state S that might obtain, you don't prefer ψ 's outcome to φ 's in S.⁹

NEVER WORSE PRINCIPLE

If φ is never worse than ψ , then you are not rationally required to prefer ψ to φ .

The motivation for this principle is somewhat similar to one given for the *Strict Dominance Principle*: if φ is never worse than ψ , then you are in a position to know that you don't prefer ψ 's actual outcome to φ 's—and if you know, as a matter of actual fact, that ψ isn't any better than φ , you aren't rationally required to prefer it.¹⁰ The case of opaque sweetening, which opened this section, serves to illustrate this constraint. In that example, because you don't prefer $A^{+\$20}$ to B and you don't prefer $B^{+\$20}$ to A, R is never worse than L^+ . And so, according to the *Never Worse Principle*, you aren't rationally required to prefer L^+ to R.

These four constraints—*Neutrality*, the *Strict Dominance Principle*, *Sensitivity to Sweetening*, and the *Never Worse Principle*—are jointly inconsistent, as the following example makes clear:

⁸ I'll mention that the analogous claim which replaces “indifference” with “ambivalence”—that is, the one that says “if you prefer X to Y and are ambivalent between Y and Z, then you should prefer X to Z”—is exactly what was being denied above when we said that ambivalence is insensitive to mild sweetening. There might be cases in which, because your preference for X over Y is so strong, you should prefer X to Z. But this principle does not hold in general—and certainly doesn't hold when your preference for X over Y is slight.

⁹ In symbols: φ is *never worse than* ψ just in case, for all state S, $(\varphi \wedge S) \not\prec (\psi \wedge S)$.

¹⁰ In the service of defending a similar principle (which they call *Competitiveness*), Bales et al. [2014] explicitly appeal to an analogy with weak dominance. Schoenfield [2014], also defending a similar principle (which she calls *LINK*, because it links facts about *expected* value to what one knows about *value*), worries that views that violate it are “imposing requirements that transcend what we actually care about: the achievement of *value*.” See [Doody, 2019b] for discussion, and critical assessment, of these arguments.

	HEADS	TAILS
L^+	A+\$20	B+\$20
L	A	B
R	B	A

Given *Neutrality*, you're rationally required to be indifferent between L and R . From the *Strict Dominance Principle*, you're rationally required to prefer L^+ to L . If you adopt the attitudes that these principles rationally require you to adopt, then, from *Sensitivity to Sweetening*, you're also rationally required to prefer L^+ to R , which violates the *Never Worse Principle*. Therefore, there can be no decision theory that both allows for ambivalence and satisfies all four of these constraints.

3 Neutrality and the Never Worse Principle

The puzzle of Opaque Sweetening has sparked a debate, which divides roughly into two camps: those who support (something like) *Neutrality* and those who support (something like) the *Never Worse Principle*.¹¹ This makes the impossibility result presented above of only limited interest. Both *Neutrality* and the *Never Worse Principle* are—as I will soon argue—needlessly contentious. The most compelling motivation for the one tells against the truth of the other (and *vice versa*), so the two aren't really *independently* plausible.

Consider *Neutrality*, for example. You should only find this principle appealing if you think that the role *states* play in determining what it's rational to do is a derivative one. The states matter only insofar as they serve to determine evaluatively relevant features of the potential outcomes and those outcomes' probabilities; once the outcomes and their probabilities have been determined, the facts about which outcome would be realized in which state is irrelevant. If two options have the same potential outcomes with the same probabilities, the reasons you have to

¹¹ In the former camp, we have: [Hare, 2010, 2013], who first introduced the puzzle and who (cautiously) endorses a principle (called "Prospects Determine Permissibility") that entails *Neutrality*; [Rabinowicz, forthcoming], who argues for a very similar principle (which he calls "Permutation (P)"); and Bader [2018]. In the latter camp, there is: [Bales et al., 2014], who defend a principle (which they call "Competitiveness") that is similar in spirit to the *Never Worse Principle*; and [Schoenfield, 2014], whose previously mentioned "LINK" principle provides a metaethical motivation for our own.

favor the one perfectly balance the reasons you have to favor the other; and so you ought to be indifferent between them.

But it's precisely this view about states that proponents of the *Never Worse Principle* should deny. On these views, the states matter (over and above their effect on the outcomes and their probabilities) because, ideally, your attitudes should match the value-relations that *actually* hold between your options—and the facts about which outcomes will be realized in which states are relevant to whether you're in a position to know which value-relation obtains in the actual world. If you know, for example, that φ 's outcome is better than ψ 's in the actual world, you should prefer φ to ψ ; if you know that you're indifferent between φ 's and ψ 's actual outcomes, you should be indifferent between φ and ψ ; and if you know—as you do with L and R —that you are ambivalent between your options' outcomes in the actual world, you should be *ambivalent*, not indifferent, between those options. Permuting states (even if they are equiprobable) might affect what you're in a position to know about how your options compare in the actual world. But, on these views, such knowledge plays an important (perhaps even *decisive*) role in determining what it's rational to do.

Just as the proponents of the *Never Worse Principle* have reason to reject *Neutrality*, the proponents of *Neutrality* have reason to reject the *Never Worse Principle*. They can claim that, while it might *often* be true that you're not required to prefer one option to another if you know that the latter is guaranteed to not be worse than the latter, this isn't *always* the case. In many cases, φ is never worse than ψ because φ 's outcomes are *at least as good* as ψ 's in each state—in which case, you can't have *more* reason to ψ than to φ . But, in this case, R is never worse than L^+ (and L^+ is never worse than R), not because R 's outcomes are at least as good as L^+ 's, but because you're *ambivalent* between their outcomes—and that's compatible with having more reason to take L^+ than R .¹² In a similar vein, they could claim that what matters in evaluating an option is not merely how good that option is relative to how things *would've* been had you chosen otherwise, but also

¹² Hare [2013, p. 51] argues along these lines. On his view, you should choose L^+ over R because you have no reason to choose R over L^+ (anything that can be said in favor of doing the former can equally well be said in favor of doing the latter) but you do have a reason to choose L^+ over R (you're sure to get twenty bucks). The fact that R is never worse than L^+ doesn't give you a reason to choose it (given that L^+ is also never worse than R), nor does it remove the reason you have to choose L^+ over R . On this view, rationality is about doing what you have the most reason to do. And, in this case, you have more reason to choose L^+ than you do to choose R .

how good things *could've* been. Were the former all that mattered, then perhaps the *Never Worse Principle* would be correct. But, because L^+ guarantees something better than what you just as well *could* get from R , you are—*contra* the *Never Worse Principle*—required to prefer L^+ to R .¹³

Both *Neutrality* and the *Never Worse Principle* are contentious. In the next section, I will show that there exists an impossibility result that follows from less controversial constraints.

4 Less Controversial Constraints

In this section, I'll develop a different impossibility result, which replaces *Neutrality* and the *Never Worse Principle* with less controversial constraints. Recall that *Neutrality* was controversial because it abstracts away from the facts about which outcomes occupy which states. Its critics contend that these facts are important because they're relevant to what you might be in a position to know about how your options compare in the actual world. In short, *Neutrality* effectively allows for cross-state comparisons—which, its critics contend, are illegitimate (at least when there's ambivalence). Also, recall that the *Never Worse Principle* was controversial because the fact that one option is never worse than another is (at least relative to one conception of reasons) nevertheless compatible with having no reason to take the former and some reason to take the latter—as illustrated by the fact that φ can be both never worse than ψ and yet stochastically dominated by it. This is, in part, because the *never worse than* -relation isn't asymmetric: it is possible for two options, like L^+ and R , to both be never worse than the other. Thus, that, e.g., R is never worse than L^+ doesn't provide you with a reason to prefer R —or, at least, not one that isn't perfectly counterbalanced by an analogous reason in the other direction. I'll introduce some constraints that avoid these criticisms.

The first is a weakening of the *Never Worse Principle*. Roughly, it says that if

¹³ Bader [2018, pg. 505] argues along these lines. On his view, you should prefer L^+ to R because the former stochastically dominates the latter. One option *stochastically dominates* another just in case, for every outcome, the probability of getting something at least as good is at least as great for the former as it is for the latter, and there is some outcome such that the probability of getting something at least as good as it is greater for the former than the latter. The probability of getting something at least as good as A is the same for both L^+ and R : 50%. Likewise for B. However, the probability of getting something at least as great as $A^{+\$20}$ is 50% for L^+ but 0% for R . Likewise for $B^{+\$20}$.

one option is never worse than another *and it is likely to be better*, then you are not rationally required to disprefer it. An option, ϕ , is *likely to be better* than another, ψ , just in case the probability you assign to the disjunction of the states in which you prefer ϕ 's outcome to ψ 's is greater than 50%.¹⁴

NEVER WORSE, LIKELY BETTER PRINCIPLE

If ϕ is never worse than ψ and is likely to be better, then you are not rationally required to prefer ψ to ϕ .

Notice that this principle, unlike the *Never Worse Principle*, makes use of a notion that is asymmetric: if ϕ is never worse and likely to be better than ψ , then ψ *isn't* never worse and likely to be better than ϕ . So, if ϕ is never worse and likely to be better than ψ , you do have a reason—a fairly powerful one, in fact—for taking ϕ over ψ (and one which isn't perfectly counterbalanced by an analogous reason in the other direction). Given this powerful consideration in ϕ 's favor, it's hard to see how you could nevertheless be rationally required to prefer ψ instead. The *Never Worse, Likely Better Principle* says that you're not. This principle is harder to deny than its cousin, the *Never Worse Principle*.

Here's an example of the principle in action. Suppose you are deciding between L^+ and R . Before you make your selection, though, the situation changes. Rather than facing a choice between the Larger box plus twenty bucks (L^+) and the Regular box (R), you are told that, either, you can pay a small fee (\$2.50, say) to acquire L^+ outright (L^\pm) or you can opt for a mixed option (M). A fair six-sided die has been rolled. If it landed on 1, you'll receive R ; if it landed on anything but 1, you'll receive L^+ . Out of these two new options— L^\pm and M —which should you choose?

Coin:	HEADS		TAILS	
	1	2–6	1	2–6
Die:				
L^\pm	$A + \$17.50$	$A + \$17.50$	$B + \$17.50$	$B + \$17.50$
M	B	$A + \$20$	A	$B + \$20$

According to the *Never Worse, Likely Better Principle*, you are not rationally required to prefer paying the \$2.50 for L^+ over letting the die decide. M , the mixed

¹⁴ In symbols: ϕ is *likely to be better* than ψ just in case $\sum_S Cr(S) \cdot 1[(\phi \wedge S) \succ (\psi \wedge S)] > 1/2$, where $1[p]$ is an indicator function that returns 1 when p is true and returns 0 when p is false. In a similar fashion, we could define notions like “*very likely to be better*” by replacing “ $1/2$ ” with a higher threshold.

option, is five-sixths likely to result in L^+ at no cost, which, no matter which prize it contains, is preferable to paying for it. Furthermore, in the unlikely event that M doesn't result in L^+ , it nevertheless promises a prize that isn't worse than the one you'd get by paying to have L^+ outright. There's certainly something to be said in M 's favor. And so any view which holds—*contra* the *Never Worse, Likely Better Principle*—that one is rationally required to prefer L^\pm to M , at the very least owes us an explanation for how that could be.¹⁵

The next constraint is best introduced with an example. Suppose that you're in possession of R —the box that contains B if the coin landed heads and A if it landed tails—and have yet to open it. Before you do, you are offered a deal: you can either accept a dollar or you can be entered into a drawing for a large sum of money. Which should you choose: the dollar or the drawing? Suppose you do some back-of-the-envelope math and discover that you are completely indifferent between the sure-thing dollar and the drawing; even though the drawing might result in a larger sum, the chance of winning it is too small for it to come out on top. Because this choice has no effect on which prize is in the box, the fact that you've yet to open the box is irrelevant to whether you should choose the dollar or the drawing; if you'd be indifferent between the two *without* R , you should be indifferent between the two *with* R as well.¹⁶

	$(1 - p)$		(p)	
<i>Drawing:</i>	LOSER		WINNER	
<i>Coin:</i>	HEADS	TAILS	HEADS	TAILS
$R^{+\$1}$	$B^{+\$1}$	$A^{+\$1}$	$B^{+\$1}$	$A^{+\$1}$
R^{+d}	B	A	$B^{+\$1/p}$	$A^{+\$1/p}$

Notice that, in comparing $R^{+\$1}$ to R^{+d} , you are in a position to know that one of the two is actually better than the other (although you know not which). Unlike the comparison between L^+ and R , where you know you are ambivalent between their respective outcomes, here, you know you're not; it could be (in the very likely event

¹⁵ For further discussion of this example, see [Doody, 2019b], which presents it as a counterexample to views, like *Prospectism* [Hare, 2010, 2013], that recommend evaluating options solely in terms of their prospects.

¹⁶ The reasoning here presupposes that you value the prizes in the outcomes *independently* of each other: that having the one does not add or subtract from the value of having the other. If that assumption doesn't hold, substitute in prizes for which it does.

that your ticket isn't a winner) that $R^{+\$1}$ is (slightly) better than R^{+d} , and it could be (in the very unlikely event that your ticket is a winner) that R^{+d} is (significantly) better than $R^{+\$1}$, but it cannot be that you're ambivalent. This is a manifestation of the fact that, although you're ambivalent between A and B, that ambivalence is orthogonal—and, thus, irrelevant—to your choice between the dollar and the drawing. As it happens, the extent to which $R^{+\$1}$ might be better than R^{+d} ($+\$1$), weighted by the chance that it *is* better ($1 - p$) perfectly balances the extent to which R^{+d} might be better than $R^{+\$1}$ ($+\$1/p$), weighted by the chance that *it* is better (p). So, because the considerations in favor of the one perfectly balance the considerations in favor of the other, you should be indifferent between the two.

In general, the thought is this. Even if you're ambivalent between some of φ 's and ψ 's outcomes, so long as you're not ambivalent between any of the outcomes *in the same state*, you should be indifferent between φ and ψ if your estimate of the extent to which φ might be better than ψ equals your estimate of the extent to which ψ might be better than φ .

ORTHOGONAL EQUIPOISE

If, in every state, there's no ambivalence between φ 's and ψ 's outcomes, then if the probabilistically-weighted average of the extent to which φ 's outcome might be better than ψ 's equals the probabilistically-weighted average of the extent to which ψ 's outcome might be better than φ 's, then you should be indifferent between φ and ψ .

This requires some further spelling out, given that, if you're ambivalent between any outcomes whatsoever, it's not straightforward how to measure the extent to which one outcome might be better than another in a way that admits probabilistic-weighting (even if you aren't ambivalent about those two outcomes specifically). There are several ways of making the notion precise. Here's one way. If your preferences could be represented with a utility-function, u , the notion is clear: the probabilistically-weighted average of the extent to which φ 's outcome might be better than ψ 's is $\sum_{S'} Cr(S') \cdot (u(\varphi \wedge S') - u(\psi \wedge S'))$, for all states S' in which $u(\varphi \wedge S') > u(\psi \wedge S')$. We can generalize the idea in the following way. Even if your preference-ordering cannot, due to ambivalence, be represented with a *single* utility-function, it can be represented with a *set* of utility-functions. Consider all of the *coherent extensions* of your preferences: that is, all of the ways of rendering

your preferences complete while holding fixed the preferences you do have.¹⁷ We then represent each of these coherent extensions with a utility-function. Let \mathcal{U} be the set of all these utility-functions. This set represents your preference-ordering. If the probabilistically-weighted averages are the same according to *all* of these utility-functions, then, so long as there's no ambivalence between φ 's and ψ 's outcomes, you should be indifferent.¹⁸

Orthogonal Equipoise is not logically weaker than *Neutrality*; the two are logically independent. But, it is, I contend, a less controversial principle in virtue of the fact that it eschews making the kind of cross-state comparisons that the opponents of *Neutrality* found objectionable. Instead, it compares the *differences* in value between the outcomes in the same states (assuming that such comparisons can be made). The constraint doesn't presume that it always will be possible to make these comparisons. But, when you can, your ambivalence is, in some sense, *orthogonal* to the choice at hand. It can, then, for the purposes of making the decision, disappear into the background.

The impossibility result, which is presented in the next section, also requires the *Weak Dominance Principle* as well as *Transitivity*.

5 The Impossibility Result

In this section, I'll show that the *Never Worse, Likely Better Principle*, *Orthogonal Equipoise*, the *Weak Dominance Principle*, *Transitivity of Preference*, *Transitivity of Indifference*, and *Sensitivity to Sweetening* are jointly inconsistent. I'll demonstrate this in a series of steps.

¹⁷ In symbols: \succeq^+ is a *coherent extension* of your preference-ordering if, and only if, (i) for all X and Y , $X \succeq^+ Y$ if you prefer X to Y or are indifferent between them, and (ii) \succeq^+ is complete.

¹⁸ This strategy makes the principle somewhat similar to Hare's *Prospectism* [Hare, 2010, 2013]. But, because the principle is restricted to cases in which there is no statewise ambivalence, *Orthogonal Equipoise* is significantly weaker. Another way of generalizing the notion takes its inspiration from Hare's *Deferentialism* [Hare, 2010]. The details are more complicated, but the rough idea is this: for each state S , we choose a representative utility-function from \mathcal{U} (subject to certain constraints), and perform the calculations using these "state-dependent" utilities. *Deferentialism* is meant to generalize the intuition underlying the *Never Worse Principle* into a full-fledged decision theory, so if the view also entails *Orthogonal Equipoise*, that's further confirmation of that principle's plausibility.

5.1 Opaque Sweetening, Redux

The first step in the argument is to show that *Orthogonal Equipoise*, along with *Dominance* and *Transitivity*, entail that, in the original Opaque Sweetening example, you are rationally required to prefer L^+ to R . The argument goes like this.¹⁹

Suppose that in front of you is a button that reads “A”. You know the following. A fair coin has been tossed. If you push the button, you win prize A if the coin landed heads and nothing if the coin landed tails. If you don’t push the button, you win prize A if the coin landed tails and nothing if it landed heads.

	HEADS	TAILS
<i>Push “A”</i>	A	o
<i>Don’t Push</i>	o	A

Because you know that the coin is fair and because you prefer prize A to nothing, *Orthogonal Equipoise* entails that you should be indifferent between pushing the button and not.

Now suppose that in front of you there is a button that reads “B”. Again, you know that a fair coin has been tossed. In this case, if you push the button, you win prize B if the coin landed tails and nothing otherwise; and if you don’t push the button, you win prize B if the coin landed heads and nothing otherwise.

	HEADS	TAILS
<i>Push “B”</i>	o	B
<i>Don’t Push</i>	B	o

For analogous reasons, *Orthogonal Equipoise* entails that you should be indifferent between pushing the button and not.

Now suppose that both buttons are in front of you. You can push neither, push one but not the other, or push both. Because, for each button, you are indifferent between pushing it and not, by *Transitivity of Indifference*, you should be indifferent between pushing neither button and pushing both.

¹⁹ This argument is discussed and presented in greater details in [Doody, 2019a].

	HEADS	TAILS
<i>Push Both</i>	A	B
<i>Push Only "A"</i>	A & B	o
<i>Push Neither</i>	B	A

Because you know the coin is fair and because (assuming you value the prizes independently) the extent to which you prefer (A & B) to B is exactly the same as the extent to which you prefer A to nothing, *Orthogonal Equipoise* entails that you should be indifferent between pushing neither and pushing only "A". For analogous reasons—this time concerning the extent to which you prefer (A & B) to A and the extent to which you prefer B to nothing—*Orthogonal Equipoise* entails that you should be indifferent between pushing only "A" and pushing both. From *Transitivity of Indifference*, then, it follows that you should be indifferent between pushing neither and pushing both.

Now, suppose we sweeten the option *Push Both* with twenty dollars.

	HEADS	TAILS
<i>Push Both</i> ⁺	A + \$20	B + \$20
<i>Push Both</i>	A	B

From the *Strict Dominance Principle*, it follows that you should prefer *Push Both*⁺ to *Push Both*. And then from *Sensitivity to Sweetening*, it follows that you are rationally required to prefer *Push Both*⁺ to *Push Neither*. Because, structurally, this decision-problem is identical to the one between *L*⁺ and *R*, this completes the first step of the argument.

For reasons that will become clear soon, notice that there's an analogous argument for the conclusion that you are rationally required to prefer *L*[±] to *R*^{+\$17}:

	HEADS	TAILS
<i>L</i> [±]	A + \$17.50	B + \$17.50
<i>R</i> ^{+\$17}	B + \$17	A + \$17

All one needs to do is to just rerun the argument from above with a \$17 bonus added to each vacation package.

5.2 The Dollar and the Drawing, Redux

The second step of the argument is to show that *Orthogonal Equipoise* and the *Strict Dominance Principle* entail that you should prefer $R^{+\$17}$ to R^* .

Coin:	HEADS		TAILS	
	1	2–6	1	2–6
$R^{+\$17}$	$B^{+\$17}$	$B^{+\$17}$	$A^{+\$17}$	$A^{+\$17}$
R^*	B	$B^{+\$20.25}$	A	$A^{+\$20.25}$

The argument is fairly straightforward. First, let's assume (as we implicitly have been) that you value money linearly, that you value the vacations and the money independently, and that you are risk-neutral. Given these assumption, you should be indifferent between a lottery that affords you a five-sixth chance of winning \$20.25 and a sure-thing \$16.875. Imagine, as we did above, that you have a choice between the two—but that, also, you are in possession of R . As we saw, *Orthogonal Equipoise* entails that, because you are indifferent between the lottery and the cash, you should be indifferent between R^* and $R^{+\$16.875}$. But, because $R^{+\$17}$ strictly dominates $R^{+\$16.875}$, the *Strict Dominance Principle* says that you are required to prefer it. And, thus, from *Sensitivity to Sweetening*, you are also required to prefer it to R^* . The extent to which $R^{+\$17}$ might be better than R^* is large enough to counteract the more likely, but less significant, extent to which the latter might be better than the former. This completes the second step of the argument: you are rationally required to prefer $R^{+\$17}$ to R^* .

5.3 The Drawing and the Roll

The third step is to argue that you are required to prefer the drawing (R^*) to the roll (M). Recall that R^* is just R sweetened with a lottery that has a five-sixth chance of paying out \$20.25, and that M is a mixture of L^+ and R , where, if you take it, you have a five-sixth chance of getting L^+ and a one-sixth chance of getting R . Let's suppose that both R^* and M turn on the same roll of the same die, as represented below.

Coin:	HEADS		TAILS	
Die:	1	2-6	1	2-6
R^*	B	$B^{+\$20.25}$	A	$A^{+\$20.25}$
M	B	$A^{+\$20}$	A	$B^{+\$20}$

The argument here is very similar to the one in the first step. First, we consider a slightly soured version of R^* that promises a potential windfall of only \$20 rather than \$20.25. Then, as in step one, we “disassemble” that choice into two different decision-problems, like so:

Coin:	HEADS		TAILS	
Die:	1	2-6	1	2-6
R_B^*	B	$B^{+\$20}$	o	o
M_B	B	o	o	$B^{+\$20}$

By *Orthogonal Equipoise*, you are required to be indifferent between R_B^* and M_B .

Coin:	HEADS		TAILS	
Die:	1	2-6	1	2-6
R_A^*	o	o	A	$A^{+\$20}$
M_A	o	$A^{+\$20}$	A	o

For the same reasons, you are required to be indifferent between R_A^* and M_A .

There are four possible combinations of actions you could take: both $M_A \& M_B$, both $R_A^* \& R_B^*$, $M_A \& R_B^*$, or $R_A^* \& M_B$. From *Orthogonal Equipoise* and *Transitivity of Indifference*, it follows that you are required to be indifferent between $(M_A \& M_B)$ and $(R_A^* \& R_B^*)$. Then, from the *Weak Dominance Principle*, it follows that you are required to prefer R^* to $(R_A^* \& R_B^*)$. From *Sensitivity to Sweetening* (and the fact that $(M_A \& M_B)$ is equivalent to M), you are required to prefer R^* to M .

5.4 The Final Step: Pay or Roll

In the first step, we reached the conclusion that you are required to prefer L^\pm to $R^{+\$17}$. In the second step, we reached the conclusion that you are required to prefer $R^{+\$17}$ to R^* . In the third step, we concluded that you are required to prefer R^* to M .

Coin: Die:	HEADS		TAILS	
	1	2–6	1	2–6
L^\pm	$A^{+\$17.50}$	$A^{+\$17.50}$	$B^{+\$17.50}$	$B^{+\$17.50}$
$R^{+\$17}$	$B^{+\$17}$	$B^{+\$17}$	$A^{+\$17}$	$A^{+\$17}$
R^*	B	$B^{+\$20.25}$	A	$A^{+\$20.25}$
M	B	$A^{+\$20}$	A	$B^{+\$20}$

From the fact that you're required to prefer L^\pm to $R^{+\$17}$ and that you're required to prefer $R^{+\$17}$ to R^* , it follows from *Transitivity of Preference* that you are required to prefer L^\pm to R^* . From that fact and that you are required to prefer R^* to M , it follows from another application of *Transitivity of Preference* that you are required to prefer L^\pm to M .

But, as we saw in the previous section, this violates the *Never Worse, Likely Better Principle*. You can be certain that, however the world turns out to be, M 's outcome won't be worse than L^\pm 's. Moreover, if the die rolls anything other than a 1, which is fairly likely to happen, L^\pm 's outcome *will* be worse than M 's. These are some powerful considerations that speak in favor of M over L^\pm , which is why the *Never Worse, Likely Better Principle* holds that you are, therefore, not required to prefer L^\pm to M . And yet, the argument in this section led us to the conclusion that you *are*. And that's a contradiction.

6 Conclusion

It's hard to choose when you're ambivalent about your choices. It's harder still—or so I hope I've shown—to choose when you are both ambivalent and uncertain. I've argued that a number of independently plausible principles governing rational choice in such situations are jointly inconsistent. I've argued that this result is, in some ways, more serious than the Opaque Sweetening puzzle because it relies on principles that are less contentious.

What should we make of this result? One possible reaction is to give up on developing a decision theory for such situations. If you're ambivalent, you should work to resolve it rather than expect decision theory to provide you with any helpful advice.

That might ultimately be correct, but I think that's to give up too early. I think that, as a matter of fact, we are all ambivalent about something at least some of the

time. I'm skeptical that one even realistically could resolve all of one's ambivalence. I'm also skeptical that one should—that it is some kind of character flaw or mark of irrationality to be ambivalent about at least some things at least some of the time. In fact, in some cases, ambivalence might be the correct—or most fitting—attitude to have, especially in the face of what might be actual value indeterminacy out in the world. And so it would be a shame—an embarrassment, even—if formal decision theory had absolutely nothing to offer us.

What are our other options? Give up one (or more) of the offending principles. I've attempted to argue that each of the principles used to derive the contradiction is fairly plausible. But plausibility is not infallibility. I think, in terms of rejecting a principle, the most promising routes are, either, to deny the *Never Worse, Likely Better Principle* or to deny *Transitivity*. Proponents of views, like *Prospectism*, which evaluate options solely in terms of the probability-distribution over its outcomes, will need to reject the *Never Worse, Likely Better Principle*. But, in doing so, they owe us an explanation for why it *seems* right—and it's not clear that the extant objections to the *Never Worse Principle* carry over. On the other hand, proponents of dominance-style principles, like the *Never Worse Principle* and the *Never Worse, Likely Better Principle*, should (I think) reject *Transitivity*. But that, too, of course comes at a cost—even a potentially non-metaphorical one (in the event that one is transformed into a money-pump).

I think more work needs to be done in exploring the terrain surrounding hard choices. I hope this chapter has helped impose some structure on that terrain. In the meantime, hard choices will remain as hard as ever.²⁰

References

R. Aumann. Utility theory without the completeness axiom. *Econometrica*, 30(3): 445–462, 1962.

²⁰ This chapter has benefited from generous contributions from many sources. I'm grateful to audiences at the Conference on Ethics and Uncertainty at the Hebrew University of Jerusalem, the departmental colloquium at the same institution, the Recent Work in Decision Theory and Epistemology workshop at Columbia University, and the Fourth Annual Meeting of the PPE Society. I'd especially like to thank Chrisoula Andreou, D Black, Luc Bovens, Jennifer Carr, Jamie Dreier, Nir Eyal, Melissa Fusco, Ze'ev Goldschmidt, Simone Gubler, Caspar Hare, Harvey Lederman, Chris Meacham, Kian Mintz-Woo, Agustín Rayo, Ittay Nissan-Rozen, Bob Stalnaker, Una Stojnić, Christian Tarsney, and likely many others. I'd also like to thank the editors of this volume, Henrik Andersson and Anders Herlitz. Any and all mistakes are my own.

- Ralf Bader. Stochastic dominance and opaque sweetening. *Australasian Journal of Philosophy*, 96(3):498–507, 2018.
- A. Bales, D. Cohen, and Toby Handfield. Decision theory of agents with incomplete preferences. *Australasian Journal of Philosophy*, 92(3):453–470, 2014.
- John Broome. Does rationality consist in responding correctly to reasons? *Journal of Moral Philosophy*, 4(3):349–374, 2007.
- Ruth Chang. Introduction. In *Incommensurability, Incomparability, and Practical Reason*. Harvard University Press, 1997.
- Ruth Chang. The possibility of parity. *Ethics*, 112(4):659–688, 2002.
- Ruth Chang. Voluntarist reasons and the sources of normativity. In D. Sobel and D. Wall, editors, *Reasons for Action*, pages 243–271. Oxford University Press, 2009.
- Ruth Chang. Hard choices. *APA Journal of Philosophy*, 3(1):1–21, 2017.
- Herman Chernoff. Rational selection of decision functions. *Econometrica*, 22:422–443, 1954.
- Donald Davidson. Hempel on explaining action. In *Essays on Actions and Events*, pages 261–275. Oxford University Press, 1976.
- Ryan Doody. Opaque sweetening and transitivity. *Australasian Journal of Philosophy*, 97(3):559–571, 2019a.
- Ryan Doody. Parity, prospects, and predominance. *Philosophical Studies*, 176(4): 1077–1095, 2019b.
- J. Dubra, F. Maccheroni, and E. Ok. Expected utility theory without the completeness axiom. *Journal of Economic Theory*, 115(1):118–133, 2004.
- 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.
- James Griffin. *Well-Being: Its Meaning, Measurement and Importance*. Clarendon Press, Oxford, 1986.
- Caspar Hare. Take the sugar. *Analysis*, 70(2):237–247, 2010.
- Caspar Hare. *The Limits of Kindness*. Oxford University Press, 2013.
- 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.
- Derek Parfit. *Reasons and Persons*. Clarendon Press, Oxford, 1984.
- Wlodek Rabinowicz. Incommensurability meets risk. unpublished, forthcoming.
- Leonard J Savage. *The Foundations of Statistics*. Wiley, 1954.
- 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.
- J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.
- Paul Weirich. *Realistic Decision Theory: Rules for Nonideal Agents in Nonideal Circumstances*. Oxford University Press, 2004.