Taking the Cake? Rational Choice in the Face of Unresolved Value Conflict

Ryan Doody

Abstract

This paper addresses a puzzle about how to choose between three options when you don't endorse a single precise way of evaluating them. It considers two different decision rules which disagree in such contexts. One is more permissive than the other. The paper argues that the best argument against the permissive rule fails, but that this argument can be rescued by making an unconventional assumption about how to model deliberation.

1 Introduction

It's irrational to select an item from a menu if that menu also includes something you strictly prefer to it.[1] Is the converse true? That is, is it rationally permissible to select an item just so long as there's nothing else on the menu you'd rather have? If yes, then your *all things considered* preferences are the ultimate arbiters of what it's rational for you to do. But in cases in which one is unable to form *all things considered* preferences over all of the items on the menu—perhaps because the choices are complex, multifaceted, and involve unresolved conflicts between competing values—it's not obvious that one's *all things considered* preferences should play such a vaunted role. Some (notably, Levi, 1986, 1999, 2006, 2008) argue that they should not: in some cases, it's irrational to choose an option even though there's nothing else on the menu you *all things consider* prefer to it. Others (notably, Sen, 1997, 2004) disagree: it's irrational to choose an option only if there's something else on the menu you'd (all things considered) rather have. Although this

[1] This claim can be resisted if the menu contains infinitely many options, none of which are maximal. Consider, for example, *Satan's Apple* (Arntzenius et al., 2004). Eve is punished if she eats infinitely many slices of the apple, but (all else equal) prefers eating more slices to fewer. No matter which slices Eve decides to eat, there'll be some way she could've done better. It's not entirely clear what Eve should do in this situation. Perhaps she is condemned to act irrationally. But, off hand, it's at least *less* irrational to eat a very, very large—but finite–number of slices than it is to eat none (and significantly *more* irrational to eat infinitely many).

paper ultimately defends the former view, I argue that, as things stand, that view is undermotivated: the best argument in its favor is fallacious. I then offer a new argument for the view, which involves complicating our picture of the role that deliberation plays in rational choice.

2 Deciding Under Value Conflict

Consider the following situation.

> **Deciding at the Diner.** You pop into a favorite local diner for an after-dinner treat. Norma brings you the diner's (unfortunately meager) dessert menu:
>
> (A) an apple pie,
>
> (B) a bowl of blueberries,
>
> (C) a cantaloupe cake
>
> When choosing a dessert, you only care about two things: *deliciousness* and *healthfulness*. You know that the apple pie is more delicious than the cantaloupe cake, which is slightly—and I stress *only slightly*—more delicious than the bowl of blueberries. You also know that the bowl of blueberries (which are chock full of anti-oxidants) is healthier than the cantaloupe cake, which is slightly—and, again, *only slightly*—healthier than the apple pie. There's no precise way in which these two considerations—deliciousness and healthfulness—weigh up against each other. They both matter to you, but aren't precisely commensurable.

Let's suppose that you don't prefer any of the items on the menu to any of the others. Your preferences are *incomplete*. Let's also suppose that we can represent deliciousness and healthfulness on a scale that allows us to assign numbers to how the options are valued relative to each of these dimensions.[2]

---

[2] The numbers are for convenience. But don't let them deceive you. The fact that the apple pie is 10 *dels* delicious and that the bowl of blueberries is 10 *healths* healthy should very much not suggest to you that the apple pie is as delicious as the bowl of blueberries is healthy. The two scales are (obviously) conventional. (We just made them up!) And we haven't said how—if at all—the two scales compare to each other. In fact, given the details of the story, it looks as though the two scales are not precisely commensurable: there's no fact about how many *healths* are worth one unit of *del*. It would be, then, just as foolish to conclude (from looking at the numbers on our scales) that the apple pie is as delicious as the bowl of blueberries are healthy as it would be to conclude that an olympic-sized pool is as long as a $50°F$ degree day is warm.

$$v_d(A) = 10 \qquad v_h(B) = 10$$
$$v_d(C) = 3 \qquad v_h(C) = 3$$
$$v_d(B) = 0 \qquad v_h(A) = 0$$

As the story suggests, the deliciousness scale, measured in *dels*, and the healthfulness scale, measured in *healths*, aren't precisely commensurable for you. There is no precise conversion rate between the two. If there were, then your decision about what to get for dessert would be, more or less, a no-brainer: convert the two amounts into some common currency and then pick the option that comes out on top. (So, for example, if each *health* were worth 2 *dels*, then you should clearly go for the blueberries.) Instead, let's consider all of the (plausible) ways the two dimensions of value could be weighed off of each other.[3] Each of these ways will correspond to a complete ranking of your options—a *coherent extension* of your incomplete preference ranking. Each is an admissible way of resolving your value conflict. Let $\mathcal{U}$, which we'll call your *representor*, be the set of all these rankings. We can use this to recover our more familiar notion of an *all things considered* preference:[4]

Preference Unanimity: You *all things considered* prefer X to Y if, and only if, for every $u \in \mathcal{U}$, $u(X) > u(Y)$.

You don't *all things considered* prefer any of the desserts to any of the others. For each pair of desserts, there are admissible ways of resolving the conflict between your values that rank the one ahead of the other and *vice versa*. Combining this view of preference with the deontic claim mentioned at the beginning—namely, that you're rationally forbidden from choosing X if, and only if, there's something on the menu you prefer to it—delivers us an attractive decision rule:[5]

---

[3] We could, instead, represent your attitude towards these options with *vectors*: e.g., $\mathbb{V}(A) = \langle 10, 0 \rangle$, $\mathbb{V}(B) = \langle 0, 10 \rangle$, $\mathbb{V}(C) = \langle 3, 3 \rangle$. And then we can articulate some principles to govern how the dimensions of value relate to *overal* value. For example, consider *Pareto*: $\langle x_1, y_1 \rangle \succ \langle x_2, y_2 \rangle$ if $x_1 > x_2$ and $y_1 > y_2$. (For a defense of this principle, see Hedden and Muñoz, 2023). But we don't accept the converse: you may prefer one thing to another even if it isn't better along *every* dimension (e.g., it can be rational to prefer $\langle 100, 9 \rangle$ to $\langle 0, 10 \rangle$). That said, we nevertheless stipulate that you lack preferences between A, B, and C.

[4] The left-to-right direction of Preference Unanimity—that if you prefer X to Y, then every function in $\mathcal{U}$ ranks X ahead of Y—is fairly uncontroversial, perhaps even true-by-definition. (This is what ensures that every $u \in \mathcal{U}$ is a coherent *extension* of your (merely partial) preference ranking.) The right-to-left direction, however, is controversial—particularly when X and Y are gambles. The controversy concerns what Hare (2010) calls cases of "opaque sweetening". (See Bader, 2018; Bales et al., 2014; Doody, 2019a,b, 2021; Schoenfield, 2014, for further discussion). Because the parties to the dispute at issue in this paper accept both directions of the principle, for presentational purposes, I will as well.

[5] Proponents of the rule (or rules similar in spirit) include Amartya Sen (2004), who calls it "intersection maximization"; Hans Herzberger (1973); and Seamus Bradley (2013). Let's say that an

Maximality: It's irrational to choose X from menu M if, and only if, there is some $Y \in M$ such that *every* admissible way of evaluating the options in $\mathcal{U}$ ranks Y ahead of X.

The rule is fairly permissive. All it takes for an action to be rationally permissible according to Maximality is for there to fail to be another option that does better relative to every admissible way of resolving the conflict between the dimensions of evaluation. In this case it says, for each of the desserts on the menu, that it's a rationally permissible choice. In particular: it's rationally permissible to choose the cantaloupe cake.

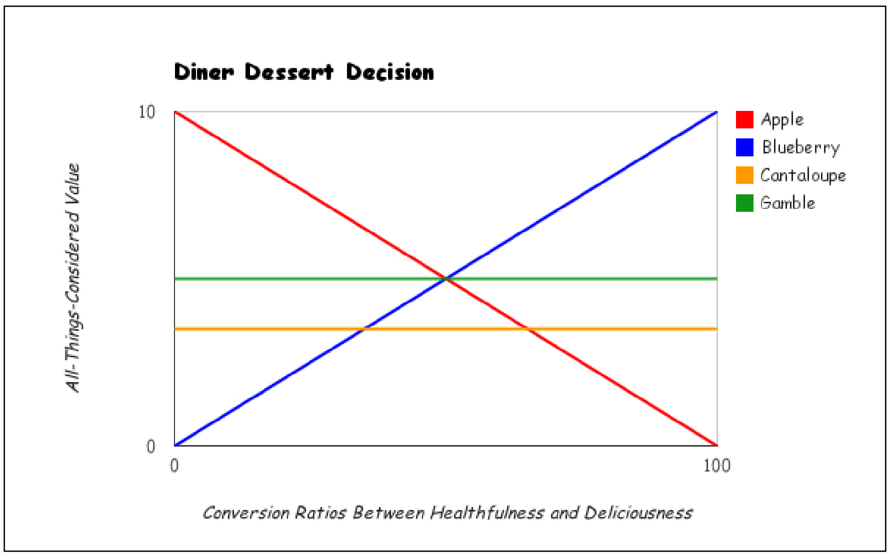

Figure 1: Values of Desserts Across Ways of Evaluations

For some, Maximality is too permissive. Isaac Levi, for example, endorses the following decision rule instead:

V-admissibility: It's irrational to choose X from menu M if, and only if, according to every admissible way of evaluating the options, there is some $Y \in M$ such that Y is ranked ahead of X.

This rule is less permissive. In order for an action to be rationally permissible, there must be some way of evaluating the options according to which it comes

option is *maximal* just in case there's nothing else on the menu you strictly prefer to it. (A maximal option needn't be *optimal*: that is, weakly preferred to every alternative.) The rule says: being maximal is necessary and sufficient for an option to be rationally permissible.

out best. In this case, there is no way of evaluating the option such that choosing the cantaloupe cake is best. (To see this, note that in order for the cantaloupe cake to come out better than the apple pie we must put a lot of weight on healthfulness, but that'll make the bowl of blueberries, which are much healthier than the cantaloupe cake, the best of the three. And, *mutatis mutandis* for deliciousness. There are some admissible ways of evaluating the options according to which the apple pie is best, and there are some admissible ways of evaluating the options according to which the bowl of blueberries is best—but there are no ways of evaluating the options according to which the cantaloupe cake is best.)

Which rule is right?

3 Levi's Argument

Isaac Levi argues that, in cases like the one presented above, there is an important distinction that Maximality cannot capture, but which V-admissibility can. In order to motivate Levi's claim, consider the following addendum to our story: Norma returns to the table to inform you that, in addition to the apple pie and bowl of blueberries but instead of the cantaloupe cake, the diner's pastry chef has made a nice carrot cake (C^+), which is only slightly less delicious than the apple pie and only slightly less healthful than the bowl of blueberries. (For definitiveness, let's say that $\mathbb{V}(C^+) = \langle 8, 8\rangle$—that is, it scores an 8 along both dimensions.) In this case, both rules agree: it's rationally permissible to choose the carrot cake.[6]

In both cases, there is a sense in which the cake (cantaloupe and carrot) comes in *second place*. But, in the first version (with C), the cake is *second worst*. While, in the second version (with C^+), the cake is *second best*. Levi argues that there *is* a sensible distinction here. And goes on to argue that because Maximality rules all three options permissible in *both* cases, it is unable to capture the distinction.

Levi says:

> [Maximality] obliterates the relevance of the distinction between case 1 and case 2 ... (Levi, 2006, 205)

And elsewhere, he says:

> [Maximality] cannot distinguish between second best and second worst cases. (Levi, 2008, 13)

To illustrate, Levi asks us to consider a slight amendment to the example. Imagine, now, that you have the option to let the decision between the apple pie (A) and

[6] There are u's in $\mathcal{U}$ that rank C^+ ahead of A, and there are u's that rank C^+ ahead of B, and there are u's that rank C^+ ahead of *both* A and B.

the bowl of blueberries (B) turn on the flip of a fair coin. If the coin lands heads, you will opt for A; if the coin lands tails, you will opt for B. Call the option to take this gamble G. Levi observes that every $u \in \mathcal{U}$ ranks G ahead of C, and ranks C^+ ahead of G.[7]

Levi points out that

> Of course, in neither case does [the agent] have the option of choosing such a lottery. But [the agent]'s attitude towards such lotteries on the counterfactual supposition that he faces a choice between the lottery and [option C] could be relevant. (Levi, 2006)

So what then is his argument? He explains,

> My argument is that such significant shared agreements [concerning the hypothetical gambles] ought to be preserved in predicaments where [the agent] is conflicted between [ways of evaluation]. (Levi, 2006)

That all might be right, but it is just a *mistake* to think that Maximality is unable to "distinguish" between the two cases, and "obliterates the relevance" of the second best/second worst distinction. Levi's own example of the hypothetical gambles brings out exactly how the proponents of Maximality *could* go about making such a distinction: in the first case, *were* the gamble an available option, then option C would be ruled out; and, in the second case, *were* the gamble an available option, then option C^+ would remain on the table. The proponents of Maximality, then, *are* able to make a distinction between the cake being second best vs being second worst—they merely disagree with Levi that the hypothetical coin-flip is relevant to what it's rational to do when the flipping the coin isn't an available alternative.

Compare this to a case of *self-binding* (viz. Arntzenius et al., 2004; Elster, 1984). If Ulysses could bind himself to keep the ship on course, there would be no need for him to *literally* bind himself to the mast. But given that Ulysses doesn't have

---

[7] This is because each $u \in \mathcal{U}$ is *expectational*: the value u assigns to a risky gamble is the weighted average of the values it assigns to the gamble's potential outcomes, where the weights correspond to those outcomes' probabilities. In particular, for any $u^* \in \mathcal{U}$, $u^*(G) = {}^1\!/\!{}_2 \cdot u^*(A) + {}^1\!/\!{}_2 \cdot u^*(B) = 5$. This way of evaluating G (where it comes out better than C and worse than C^+) can be resisted—either by denying (as mentioned in footnote 4) the right-to-left direction of Preference Unanimity, or by denying that the u's in $\mathcal{U}$ should be expectational, or by rejecting the entire "sets of utility functions" framework. Furthermore, there is a good, principled reason for denying that G should be valued in this way: you don't prefer either of its outcomes (A or B) to C, and you don't prefer C^+ to either of its outcomes. Are you really rationally required to prefer the gamble to C when you know that there's no way for the gamble to result in something you prefer to C? (See Bales et al., 2014; Doody, 2019b; Hare, 2010; Schoenfield, 2014, for arguments that you aren't required to prefer the gamble). But there are also compelling reasons to think that gambles must be valued in this way.

the option to ignore the call of the Sirens through a sheer act of will, it would be foolish to make his decision about what to do based on what would best in the hypothetical situation in which he could. Similarly, proponents of Maximality should say that, while it might be irrational to choose C when you have the option to take the gamble instead, it needn't be irrational to choose C when the gamble isn't on the proverbial menu.

In short, then, if Levi's complaint is that Maximality cannot make the same distinctions that V-admissibility can, he is simply mistaken. The difference, for proponents of Maximality, between C and C^+ is that it *would* be irrational to choose C but not irrational to choose C^+ *were* you able to flip a coin instead. On the other hand, if Levi's complaint is that Maximality cannot make the same distinctions *in the same way* that V-admissibility can, then he's correct—but that clearly begs the question.

4 Deliberation and Agenda-sensitivity

As things stand, Levi's argument against Maximality fails. But his observation that it *would* be irrational to choose C if a 50/50 gamble between A and B were also on the menu is suggestive. It points us toward a more compelling (but, as I will soon argue, also unsuccessful) objection. Here's the rough idea.

Imagine approaching your decision in the following way: you, first, are going to decide whether or not to choose C; if you decide to reject C, then you'll go on to decide between the two remaining options: A and B (Figure 2). This seems like a perfectly fine way to make a decision between three options. And, given that you know that you lack a preference between A and B, it's not unreasonable to be uncertain about which you would choose were you to face a choice between the two. However, if you are genuinely uncertain about what you would choose were you to reject C in favor of deciding between A and B, you might think of your future-self—at least from your current perspective—as akin to the flip of a coin. You have no good reason to think you're any more likely to select the one over the other.[8] In which case, if this is how you are approaching the decision, you are (in some sense) facing a choice between C and a 50/50 gamble between A and B. But, as mentioned in §3, every $u \in \mathcal{U}$ ranks G (the 50/50 gamble between A and B) ahead of C. It follows, from Preference Unanimity, that you *all things*

[8] Because you lack a preference between A and B, you have no rational basis for *choosing* one over the other. This is—to borrow a distinction from Ullmann-Margalit and Morgenbesser (1977)—a situation that calls for *picking* rather than *choosing*. When we choose one option over another, we do so on the basis of the reasons we have that favor the one over the other. One way to think about picking, on the other hand, is that "when we are in a genuine picking situation we are in a sense transformed into a chance device that functions at random and effects arbitrary selections" (Ullmann-Margalit and Morgenbesser, 1977, p. 773).

Reject the Cake?

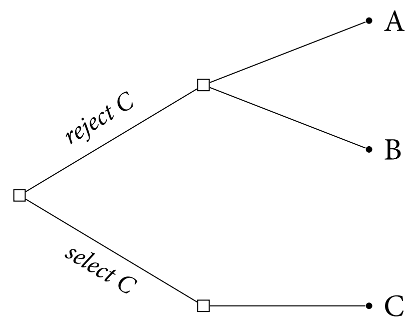

Figure 2: Your choice to either select C or to decide between A and B

considered prefer G to C. And so, C is not maximal—there's something you prefer to it—and, thus, its not rational to choose it.

According to MAXIMALITY, what it's permissible for you to choose will depend on the way in which you structure your deliberation. It's permissible to choose C …unless you are uncertain about what you would choose after rejecting C, in which case, choosing C is irrational. V-ADMISSIBILITY, on the other hand, isn't sensitive to how you structure the choice between the three options. It's irrational to choose C either way.

But now here is a plausible thought. What it's rational for you to do should be sensitive *only* to what you care about and to how you take the world to be. It shouldn't be sensitive to how you structure your deliberation between your options. If you're rational, how the decision is *framed*—how you set the *agenda*—shouldn't matter.[9] But, for MAXIMALITY, it very well might. And that's objectionable.

4.1 Unpacking the Objection

That's the objection. It relied on several claims. Let me unpack them in more detail.

P1 In choosing an option from $\{A, B, C\}$, it's rationally permissible to set the agenda as it's represented in Figure 2.

[9] In a series of influential papers (Tversky and Kahneman, 1979, 1984, 1986), Daniel Khaneman and Amos Tversky argue that, for many of us, how a decision-problem is framed *does* matter. However, their examples pertain to how particular *options* are framed—that is, how various features of an option are described—and not to the way in which those options are considered when deliberating. Holding fixed how the options themselves are described, should it matter e.g. the order in which they are considered? Furthermore, Khaneman and Tversky are engaged in a description project, not a normative one. They agree that principles that forbid framing effects are "normatively essential" (Tversky and Kahneman, 1986).

Although we will consider rejecting this claim, off hand, it appears plausible enough. There are various ways to decide between multiple options, and this way—starting with the question of whether to select or to reject C—seems as good as any other. Furthermore, notice that if one of your options is optimal (i.e., weakly preferred to the others), this way of deciding would guarantee its selection. (And, if several of your options are optimal, you are guaranteed to select one of them.) Relatedly, this method guarantees that you won't select an option if you disprefer to some other. And, while it needn't be irrational to non-instrumentally value some procedures over others, let's assume that you value *ways of deciding* only instrumentally—in terms of the values you assign to the options that might result from employing it.

P2 In facing a choice from $\{A, B\}$, it's rational to assign equal credence to ending up with either.

You know that you lack a preference between A and B. And so you have no more reason to think, if you'll be faced with a choice between them, that you'll select A over B than that you'll select B over A. And so it's rational to regard either possibility as equally likely.[10]

P3 If you assign equal credence to ending up with either A or B, the prospect of selecting an option from $\{A, B\}$ corresponds to a 50/50 gamble between A and B.

Like before, let's assume that you don't derive any additional value from deciding from $\{A, B\}$ that isn't exhausted by the value of whichever option is ultimately

[10] The reasoning here appears to appeal to something like the *Principle of Indifference*, which is controversial (see, for example, Hajek, 2003, p. 187-188). I don't think the argument turns on the truth of this principle in its full generality, however. For our purposes, it's enough that it be reasonable for you to assign credences to your future actions in the way described above. It needn't be the case that you *must*—on the pains of irrationality—do so, only that this is an epistemically reasonable reaction to your situation. Given that you have no more reason to think you'll pick A over B, it isn't unreasonable for you to distribute your credences uniformly.

That said, one could argue that, in a case like this, (i) you are *radically uncertain* about what you might go on to pick, and (ii) the uniquely epistemically rational response to radical uncertainty is to adopt *imprecise* probabilities over the various possibilities. (Seidenfeld (1988, p. 310-311), discussing a closely related problem, advocates representing "this uncertainty with a (maximal) convex set of personal probabilities" over the admissible options.) Generalizing a decision theory to handle imprecise probabilities (in addition to unresolved value conflict) is no easy task—but it's certainly not implausible. If the resulting decision theory is sufficiently permissive (for example, if it says that an option is permissible just in case there is some probability and utility functions in your representors according to which that option maximizes expected utility), you would not be required to reject C. Consider, for example, the probability function that's heavily biased towards selecting A, and the utility function that's heavily biased towards healthfulness. Depending on the details, the uncertain prospect of selecting from $\{A, B\}$ can have lower expected utility than C, according to those functions. This isn't terribly surprising given that a permissive, doubly-imprecise view is extremely capacious. Exploring views of this kind in more detail is beyond the scope of this paper.

selected. The value to you, then, of the prospect of selecting an option from $\{A, B\}$ is a function of how you value the two possible outcomes—getting A or getting B—weighted by how likely you take those outcomes to result. From **P2**, you assign equal credence—50/50—to both possibilities. And so the prospects of going on to select from $\{A, B\}$ corresponds to a 50/50 gamble between A and B.

P4 Every $u \in \mathcal{U}$ ranks G ahead of C.

As mentioned in footnote 7, this is controversial. Assuming that every $u \in \mathcal{U}$ is expectational, though, it's easy to see that this is true. (What is controversial, of course, is whether we should require every u to be expectational.) Let $u^* \in \mathcal{U}$ be an arbitrary (admissible) way for the two dimensions of value to be weighed off against each other. And let's say that, in general, $u^*(\langle x_1, y_1 \rangle) = \alpha \cdot (x_1) + (1-\alpha) \cdot (y_1)$, where $\alpha : 1 - \alpha$ is a conversion ratio between *dels* and *healths*.[11] Because these dimensions aren't precisely commensurable for you, there are many conversion ratios compatible with your attitudes—and each $u \in \mathcal{U}$ corresponds to a way of converting the two dimensions into a common currency.

$$\begin{aligned} u^*(G) &= {}^1\!/_2 \cdot u^*(A) + {}^1\!/_2 \cdot u^*(B) \\ &= {}^1\!/_2 \cdot \big(\alpha \cdot (10) + (1-\alpha) \cdot (0)\big) + {}^1\!/_2 \cdot \big(\alpha \cdot (0) + (1-\alpha) \cdot (10)\big) \\ &= \alpha \cdot (5) + (1-\alpha) \cdot (5) \\ &= 5 \end{aligned}$$

$$\begin{aligned} u^*(C) &= \alpha \cdot (3) + (1-\alpha) \cdot (3) \\ &= 3 \end{aligned}$$

$u^*(G) > u^*(C)$. And because u^* was chosen arbitrarily, the inequality holds for every $u \in \mathcal{U}$.

P5 If every $u \in \mathcal{U}$ ranks G ahead of C, you *all things considered* prefer G to C.

As previously mentioned, this is the (controversial) right-to-left direction of Preference Unanimity. Together, **P4** and **P5** say that you should prefer G to C. But

[11] Taking u^* to have this form might, understandably, seem undermotivated. Why, for example, are we taking the weighted average of these values, as opposed to performing any other mathematical operation? This is a fair question. But there are things to say in response. Although I will not do so here, it's possible to provide a justification for why these functions should take this form. Very briefly, if your preferences over the value vectors obey certain constraints—e.g., *Pareto*, *Indifference Pareto* (i.e., $\langle x_1, y_1 \rangle \approx \langle x_2, y_2 \rangle$ if $x_1 = x_2$ and $y_1 = y_2$), are expectational, etc.—we can prove that the u's must be additive. (The result exploits the structural similarity between aggregating dimensions of value into an *overall* ranking and aggregating individual utilities into a social welfare function. In particular, the axiomatic "utilitarian theorem" of Harsanyi, 1955). Whether these are plausible constraints or not is another matter.

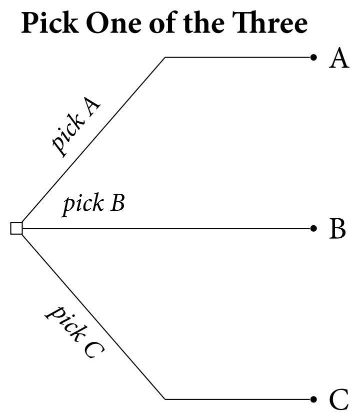

Figure 3: Your choice to pick one of the three

you know, no matter how G resolves, you don't prefer G's outcome to C. It's (at the very least) quite strange to prefer the gamble and yet know that you don't prefer its outcome. There's much that can be said about this. But now is not the time.

P6 If there's an option you prefer to X, it's not rationally permissible to choose X.

This is something both Maximality and V-Admissibility can agree on. It's irrational to take an option when you could instead take something you'd rather have.[12]

From these seven claims it follows that, in deciding between $\{A, B, C\}$, (i) it's permissible to decide by setting the agenda as depicted in Figure 2, but (ii) if you do, it's irrational to choose C. However, there are other ways of structuring your deliberation between the desserts—other ways to set the agenda—according to which it *is* rational to choose C. For illustration, here are three.

(1) **Pick One of the Three.** Knowing that you lack an *all things considered* preference between all three, you could elect to simply pick one of the desserts at random (Figure 3). Because C isn't dispreferred to either A or B, according to Maximality, it's rational to go ahead and pick it.

(2) **Reject the Pie?** Or, instead, you could turn your focus to A: first, decide whether or not to select it; and, if you elect to reject it, go on to decide between the remaining two options: B or C (Figure 4). Because you lack a

[12] Echoing what was said in footnote 1, understand **P6** to be restricted to menus in which there exists at least one maximal item. If the menu contains infinitely many items, there might fail to be a maximal element (for the same reason there fails to be a highest natural number). And it's unclear what, if anything, it's rational to do when faced with such a menu. (Thankfully—or, I suppose, not—we are unlikely to face such a menu.)

Reject the Pie?

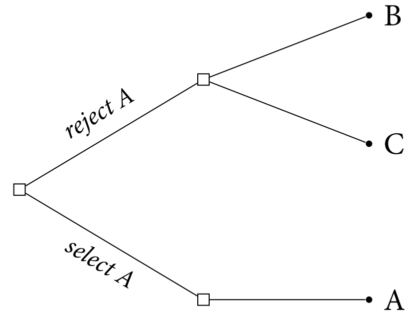

Figure 4: Your choice to either select A or to decide between B and C

preference between B and C, it's not unreasonable to be uncertain about which you would choose if given the choice. And because you don't have good reason to be more confident that you'd select one rather than the other, going on to decide from $\{B, C\}$ corresponds to a 50/50 gamble between the two. Because both B and C score more highly along the healthfulness dimension than A, there are some $u \in \mathcal{U}$ that rank the gamble ahead of A. And because A scores more highly along the deliciousness dimension than both B and C, there are also some $u \in \mathcal{U}$ that rank A ahead of the gamble. So, you lack a preference between selecting and rejecting A, which (according to Maximality) means that it's rationally permissible to reject A. And, because you don't prefer B to C, it's rational to—after having rejected A—go on to select C.

(3) **Dessert Tournament (C gets a Bye).** Or, your could focus on two of the desserts (e.g., A and B), decide which out of the two to select, and then have the winner face-off against the remaining option (e.g. C)—a *dessert tournament* of sorts (Figure 5). Because you lack a preference between A and C, you might be uncertain about which you'd end up with if faced with the choice between the two. Likewise for B and C. And so, selecting A over B corresponds to a 50/50 gamble between A and C, whereas selecting B over A corresponds to a 50/50 gamble between B and C. You don't prefer either gamble to the other—there are some us in $\mathcal{U}$ that rank the one ahead of the other, and there are other us in $\mathcal{U}$ that rank the gambles the opposite way. And so, according to Maximality, it's rational to select either—and then it's rational to go on to select C.

The point is that there are various rational ways to decide between $\{A, B, C\}$. And if Maximality is correct, these various ways of structuring your deliberation

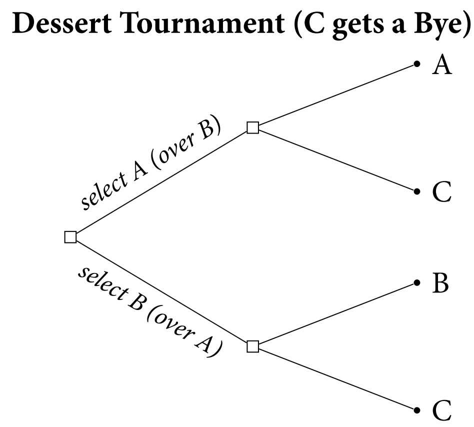

Figure 5: You choose between A and B, the winner competes against C in the next round

disagree about whether it's permissible to select C. And so, according to Maximality, whether it's rational for you to select C depends on how you frame the decision. But—so the objection goes—what it's rational to do shouldn't be sensitive to such things.

Features about how you structure your deliberation (e.g., the order in which you consider your options) very well might influence what you end up doing, but such things shouldn't affect what it's *rational* for you to choose. Whether it's irrational to take the cake should depend on features *of the cake*—its deliciousness, its healthfulness, etc.—not on contingent facts about how you've elected to deliberate, like the order in which you're considering the options. But Maximality, implausibly, makes what it's rational to do sensitive to such things. And *that* takes the cake.

4.2 Response 1: Not All Agendas are Rational

As we've seen, when a menu contains more than two items, there are various ways to go about selecting from it. We looked at a few different ways to set one's deliberative agenda when selecting an option from the menu $\{A, B, C\}$. And we saw that, according to Maximality, it's rational to choose C relative to some deliberative agendas (e.g., **Pick One of the Three**), but not others (e.g., **Reject the Cake?**). And that, for this reason, Maximality is objectionably sensitive to features that it shouldn't be.

But, one might rightly worry, that the rule's sensitivity to which deliberative agenda is employed is only objectionable if it is *rational* to set the agenda in those ways. (There's nothing incoherent about issuing recommendations for what to do if you're in a situation that you rationally shouldn't be in.) But if every deliberative agenda on which, according to Maximality, it is irrational to select C is one that it wouldn't be rationally permissible to employ anyway, it's much less objectionable that the rule issues verdicts different from those it offers relative to the agendas it *is* rational to employ. The objection to Maximality implicitly relied on the thought that the various deliberative agendas were on equal footing.

Whether it's rational to set one's deliberative agenda a particular way depends on how valuable it would be to set the agenda that way, what other ways of setting the agenda there are, and how valuable it would be to set the agenda those ways instead. As before, we will assume that ultimately you only care about the desserts—you don't value the anything *intrinsic* about the deliberative agendas themselves. And so, like before, the value you place on deciding in some particular way will be a function of the values you assign to the desserts and how likely you think it is that you'll end up with one thing rather than another.

For illustration, let's revisit the agendas from earlier.

(1) **Pick One of the Three** (Figure 3). Because you lack a preference between all three desserts, you have no reason to think that, if you were to go ahead and pick one of the three, you'll be any more likely to end up with one of the prizes than the others. And so, it's not unreasonable to split your credence equally between all three possibilities: that you pick A, that you pick B, and that you pick C. So, this way of setting the agenda corresponds to a gamble that pays out each prize with equal probability: $1/3$. And so, the value of deciding in this way corresponds to $\langle 4^1/_3, 4^1/_3 \rangle$.

(2) **Reject the Pie?** (Figure 4). As we saw before, rejecting A corresponds to a 50/50 gamble between B and C. It's value, then, is $\langle 1^1/_2, 6^1/_2 \rangle$. The value of selecting A is $\langle 10, 0 \rangle$. You don't, then, prefer selecting A to rejecting it—and so, it's permissible to do either. It's not unreasonable to be uncertain about what you would do in this situation—taking both possibilities to be equally likely. And so, the prospect of facing this way of deciding corresponds to a 50/50 gamble between $\langle 1^1/_2, 6^1/_2 \rangle$ and $\langle 10, 0 \rangle$, which corresponds to $\langle 5^3/_4, 3^1/_4 \rangle$.

(3) **Dessert Tournament (C gets a Bye)** (Figure 5). As before, selecting A over B corresponds to a 50/50 gamble between A and C, and selecting B over A corresponds to a 50/50 gamble between B and C. The former corresponds to $\langle 6^1/_2, 1^1/_2 \rangle$. The latter corresponds to $\langle 1^1/_2, 6^1/_2 \rangle$. You don't prefer one to the other, so it's permissible to select A over B and it's permissible to select B over A. And so it's not unreasonable to be uncertain about which of the two you'd choose if faced with the choice—giving equal probability to each possibility. As a result, the prospect of deciding in this way corresponds to a 50/50 gamble between $\langle 6^1/_2, 1^1/_2 \rangle$ and $\langle 1^1/_2, 6^1/_2 \rangle$, which comes out to $\langle 4, 4 \rangle$.

(4) **Reject the Cake?** (Figure 2). Selecting C is valued at $\langle 3, 3\rangle$. On the other hand, rejecting C corresponds to a 50/50 gamble between A and B, which is valued at $\langle 5, 5\rangle$. You prefer the latter to the former, so it's not permissible to select C. The prospect of deciding this way, then, is valued at $\langle 5, 5\rangle$.

Faced with a choice over, not $\{A, B, C\}$, but the various ways of setting the agenda for deciding amongst them—a meta-decision—there are some agendas you prefer to others. In particular, you prefer the **Reject the Cake? Agenda** ($\langle 5, 5\rangle$) to the **Pick One of the Three Agenda** ($\langle 4^1/_3, 4^1/_3\rangle$) to the **Dessert Tournament (C Gets a Bye) Agenda** ($\langle 4, 4\rangle$). It's irrational to choose something when you could instead have something you prefer. So, it's irrational to employ **Dessert Tournament (C Gets a Bye)** and **Pick One of the Three**. But neither **Reject the Pie?** nor **Reject the Cake?** are preferred to the other—so both of these ways of deciding are rational for you to employ. But now notice the following: with respect to the former agenda, it's rational to select C (because it's permissible to reject A and it's permissible to select C over B); but, with respect to the latter agenda, it is *ir*rational to select C (because, if you're uncertain about which of A or B you'd select, you prefer rejecting C to selecting it).

This blocks the response to the objection to Maximality. There are (at least) two ways to set the deliberative agenda that are rational, and that disagree about whether it's rational to select C. Some of the deliberative agendas did end up being irrational—they were worse along both dimensions than other ways of setting the agenda. But there is at least one agenda that it's permissible to employ and that according to which it is irrational to select C.[13]

Furthermore, notice that *all* of these deliberative agendas (even the ones that are irrational to employ) are preferred to selecting C outright ($\langle 3, 3\rangle$). There are some rational ways of deciding relative to which it is *rational* to select C, there are some rational ways of deciding relative to which it is *irrational* to select C, and that you prefer employing any of these ways of deciding to selecting C directly.

4.3 Response 2: "*Tu quoque*, V-Admissibility"

The objection to Maximality was, roughly, that whether it's rational or irrational for you to select C depends on how you've set the deliberative agenda. But what it's rational to do shouldn't be sensitive to such things.

[13] We've only looked at four ways of setting the deliberative agenda, but there are several more. They don't undermine the claims made in this paragraph, though, because **Reject the Cake?** continues to be maximal. All of the remaining ways to set the agenda either correspond to $\langle 3^1/_4, 5^3/_4\rangle$ or $\langle 5^3/_4, 3^1/_4\rangle$. The value of the version of **Dessert Tournament (A gets a Bye)** corresponds to the latter, whereas the value of the version in which **B gets a Bye** and the value of **Reject the Bowl?** both correspond to the former.

You might find this to be a convincing standalone objection against Maximality. However, in order for it to serve the grander dialectical purpose of providing support for V-Admissibility over Maximality, it would need to be the case that V-Admissibility's recommendations aren't sensitive to how you've set the deliberative agenda. Unfortunately, this is not the case. *Both* rules fall prey to the objection. And that should make you question whether the objection succeeds in either case. At the very least, the objection—insofar as it undermines both rules—fails to support one over the other. We should either somehow reject the objection, reject both of the views, or revise the objection so that it no longer applies to both. At the end, I will briefly sketch a new version of the objection that applies to Maximality alone.

First, allow me demonstrate how the current version of the objection undermines V-Admissibility as well. According to that view, it's not rational to choose C from $\{A, B, C\}$. As we've just seen, there are ways to decide between those options—e.g., **Reject the Pie?** —according to which it *is* rational to select C. However, it would be too hasty to immediately conclude from this that, for V-Admissibility, whether it's rational for you to select C depends on how you've set the deliberative agenda. The reason (which should be familiar from the previous section) is that the ways of setting the agenda that rationalize selecting C (e.g., **Reject the Pie?**, etc.) might not themselves be rational to employ. They might be, given all the other ways you might decide between the desserts, irrational.

In the previous section, we illustrated that some of the agendas that rationalized selecting C were themselves rational to employ. However, in that argument, the sense in which those agendas were "rational" is that they satisfied Maximality: there weren't other agendas that you *all things considered* preferred—i.e., there was nothing that was better along all dimensions of value. But, recall, V-Admissibility employs a different (and less permissive) standard of rationality: there must be some $u \in \mathcal{U}$ that ranks it on top. But, given that one of the things you could do is select A outright (the value of which is $\langle 10, 0 \rangle$) and another thing you could do is select B outright (the value of which is $\langle 0, 10 \rangle$), it won't be rational—by the lights of V-Admissibility—to employ an agenda like **Reject the Pie?**, whose value is $\langle 5^3/_4, 3^1/_4 \rangle$. This is because, while it's true that you don't prefer some other way of making the decision, there nevertheless is no $u \in \mathcal{U}$ that ranks **Reject the Pie?** on top.[14] And the same goes for all of the other agendas ac-

[14] To see this, notice that in order for a $u \in \mathcal{U}$ to rank that agenda ahead of *both* selecting A outright and selecting B outright, there must be an $\alpha \in [0, 1]$ such that:

$$\alpha \cdot (10) + (1 - \alpha) \cdot (0) < \alpha \cdot (5^3/_4) + (1 - \alpha) \cdot (3^1/_4)$$
$$\alpha < 0.4\bar{3} \text{ and}$$
$$\alpha \cdot (0) + (1 - \alpha) \cdot (10) < \alpha \cdot (5^3/_4) + (1 - \alpha) \cdot (3^1/_4)$$
$$0.54 < \alpha$$

Reject the Cake? (Carrot version)

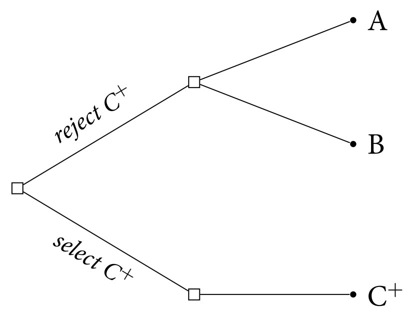

Figure 6: You decide whether to select or reject the carrot cake

cording to which it's rational to end up selecting C.[15] V-Admissibility says that it's irrational to choose C from $\{A, B, C\}$, and that it's irrational to set the agenda in a way that would make it rationally permissible to end up with C. And so, for V-Admissibility, whether it's rational to select C from $\{A, B, C\}$ *doesn't* depend on how you think about it.

But V-Admissibility is not yet out of the woods. There are other decision problems in which what it's rational to do (according to V-Admissibility) *does* depend on how you think about it. Consider an example from before: choosing a dessert from the menu $\{A, B, C^+\}$, where C^+ is a delicious and healthful carrot cake, scoring an 8 along both dimensions ($\langle 8, 8\rangle$). And consider again the **Reject the Cake?** agenda, but where now the cake is carrot, not cantaloupe (Figure 6).

As before, we can think of rejecting C^+ as a 50/50 gamble between A and B, which has value $\langle 5, 5\rangle$. Because the value of C^+ is $\langle 8, 8\rangle$, selecting C^+ is not only rational, it's *rationally required*. But V-Admissibility says, when choosing from $\{A, B, C^+\}$, it's permissible to take any of the options. And so, according to V-Admissibility, selecting C^+ is rational but not required—it would be okay to select, e.g., A instead.

And that means that V-Admissibility, too, entails that whether it's rational to do something (e.g., select A) depends on how you think about it. It's rationally

---

But no number is both greater than 0.54 and less than $0.4\overline{3}$.

[15] In particular, every $u \in \mathcal{U}$ ranks **Reject the Cake?** ahead of **Dessert Tournament (C Gets a Bye)** and ahead of **Pick One of the Three**, and neither of those two are rational ways to make the decision. Of the remaining agendas, **Reject the Pie?** and **Dessert Tournament (A Gets a Bye)** both have value $\langle 5\tfrac{3}{4}, 3\tfrac{1}{4}\rangle$. And so, as footnote 14 illustrates, for every $u \in \mathcal{U}$, there'll be something else you could do—either select A outright or select B outright—that is ranked higher. And so—at least according to V-Admissibility, those two aren't rational ways to make the decision either. Finally, **Reject the Bowl?** and **Dessert Tournament (B Gets a Bye)** both have value $\langle 3\tfrac{1}{4}, 5\tfrac{3}{4}\rangle$. And so, for reasons symmetric to those above, those two aren't rational ways to make the decision either. Aside from selecting one of the options outright, the only remaining agenda—**Reject the Cake?**—*is* a rational way to make the decision (even by V-Admissibility's standards). But, on that way of deciding, it's irrational to end up with C.

permissible to select A from the menu *if you do so outright*. However, if you deliberate by, first, considering whether or not to take the cake (and, in doing so, remain uncertain about whether you'd select A or B if you have the choice), it's *not* rationally permissible to select A. It's irrational to not take the cake.

5 Endorsable Choice

The previous section concerned a particular normative standard for rational choice: *agenda-insensitivity*. According to that standard, a choice can be rational only if it's a permissible choice to make in *all* of the rationally permissible ways of making that decision. As we saw, that standard is arguably too strong.

In this section, I will do two things. First, I will draw out a consequence of accepting *agenda-insensitivity*, and respond the objection that the standard is *so* strong that it makes rational choice impossible. Second, I will sketch a weaker (but still compelling) normative standard, and show that it supports V-Admissibility over Maximality.

Let me introduce some notation. Let M be a finite menu of options. And let M^* be the menu consisting, intuitively, of all of the ways of deciding between the options in M. More precisely, M^* is a menu of all the menus that can be made from the members of M, the menus that can be made from those menus, and so on and so forth. Let $\mathsf{c}(M) \subseteq M$ be M's *choice set*—i.e., the set consisting of those members of M that it is rationally permissible to choose. Because we're looking at menus consisting of other menus (which might contain other menus, and so forth), let $\mathsf{c}^\star$ be the transitive closure of the choice-function c, defined as follows: $X \in \mathsf{c}^\star(M)$ if and only if $X \in \mathsf{c}(M)$, or there's some $Y \in \mathsf{c}(M)$ such that $X \in \mathsf{c}(Y)$, or there's some $Z \in \mathsf{c}(M)$ and some $Y \in \mathsf{c}(Z)$ such that $X \in \mathsf{c}(Y)$, or …and so on.

> *Agenda-insensitivity*:
>
> $X \in \mathsf{c}(M)$ only if, for *every* $Y \in \mathsf{c}(M^*)$, $X \in \mathsf{c}^\star(Y)$.

This standard places a heavy constraint on rational choice. It says: it's rationally permissible to choose X from menu M only if, for every way of deciding between the items on M, it is rationally permissible to choose to decide in that way and, if you did choose to decide in that way, you could make a series of rationally permissible choices that will ultimately result in X. This standard is motivated by the idea that whether it's rational to do something or not shouldn't depend on the deliberative procedure that's used (at least so long as that deliberative procedure is itself something it's rational to employ).

Although there is something intuitively compelling about that thought, it might be objected that *Agenda-insensitivity* is much too strong. For example, it appears to entail that, when choosing from $M = \{A, B, C^+\}$, there's nothing it's rationally permissible to choose. It's not rational to choose A because there's a rationally permissible way to make the decision (see Figure 6) in which A isn't a rational choice. It's not rational to choose B for the same reason. But it is also not rational to choose C$^+$. One way to make a decision between $M = \{A, B, C^+\}$ is to choose A outright. This way of making the decision—choosing A—is both maximal (there's nothing on the enriched menu, M^*, that's preferred to it) and V-admissible. But, if that's a rationally permissible way to make your decision, none of the items on the menu are rationally permissible—because C$^+$ can't possibly result from choosing A outright. And so none of the options are a rational result from *every* rational way of making the decision. And so, $\mathsf{c}(\{A, B, C^+\}) = \emptyset$.

This objection is a bit too quick because it's not obvious that selecting $\{A\}$ from M^*, which corresponds to choosing A outright, *is* rationally permissible. It's true that there's nothing else on M^* that you strictly prefer to it (and that it's V-admissible, etc.), but those are only necessary conditions for it to be rational to choose $\{A\}$ from M^*. Furthermore, if we accept *Agenda-insensitivity*, it follows from the argument in the previous paragraph that $A \notin \mathsf{c}(\{A, B, C^+\})$. Surely an analogous argument can be used to show that $\{A\} \notin \mathsf{c}(M^*)$.

Requiring an option to be a rational choice from *every* rational way of making the decision is much too high of a bar. But, noticing that, suggests a weaker normative standard that, nevertheless, captures much of what was compelling about *Agenda-insensitivity*. Here's the idea. Instead of requiring that an option be rationally achievable from *every* rational way of making the choice, only require the option to be rationally achievable from *some* rational way of making the choice. As we saw, there are some ways of making a decision in which an option isn't rationally achievable simply because it isn't achievable at all. But, even when it is, why think that how you structure your deliberations should have *no* effect on what it's rationally permissible for you to choose. Instead, we might only require that an option be *endorsable*: that there be some rationally permissible way of making the decision such that, were you to make the decision that way, you could make a series of rationally permissible choices that would result in that option.

> *Endorsability*:
>
> $X \in \mathsf{c}(M)$ only if, there is some $Y \in \mathsf{c}(M^*)$, such that $X \in \mathsf{c}^\star(Y)$.

Furthermore, accepting *Endorsability* requires us to reject Maximality. Recall the choice between the original three desserts $\{A, B, C\}$. The enriched menu of all of the ways to choose between the three desserts contains ways of deciding that have values corresponding to a 50/50 gamble between A and B ($\langle 5, 5\rangle$). Because you prefer ways of deciding that have this value to C ($\langle 3, 3\rangle$), and because the menu contains *all* of the ways to decide between the options—including higher-and higher- order ways of deciding between the ways of deciding between the desserts—none of the options in which C is rationally achievable are maximal. The reason is that, if C is rationally achievable from some way of deciding, the value of that option reflects the fact that, if you select it, C has some probability of resulting. But because you disprefer C to other items on the menu, there will be other ways of making the decision—ways in which C is *not* rationally achievable—to which you assign higher value. So, C is not *endorsable*. But MAXIMALITY holds it to be permissible nonetheless.

Accepting *Endorsability*, however, at least as far as I can tell, doesn't require us to reject V-ADMISSIBILITY. Recall the choice from $\{A, B, C^+\}$. Because the carrot cake is so attractive, there are rationally permissible ways of making the decision such that, if you make the decision that way, you are rationally *required* to choose it. As a consequence, A and B were not rationally achievable from *every* rational way of making the decision. But *Endorsability* doesn't require an option to be rationally achievable from *every* rational way of making the decision, only from *some*. And, in this case at least, all three of the options are endorsed by at least one rationally permissible way of making the decision.

6 Conclusion

V-ADMISSIBILITY say, when facing the menu $\{A, B, C^+\}$, it's permissible to take any; but, when facing the menu $\{A, B, C\}$, only A and B are permissible. *Endorsability* says that a choice is rational only if there is some rationally permissible way of deciding between your options that endorses that choice. In the choice between $\{A, B, C^+\}$, all three of the options are endorsable; but, in the choice between $\{A, B, C\}$, only A and B are endorsable. MAXIMALITY, on the other hand, holds that all three options on both menus are rational. In particular, MAXIMALITY says that it is rationally permissible to choose C from $\{A, B, C\}$—because you don't prefer A to C and you don't prefer B to C. But there is no rationally permissible way of making the decision that recommends C. It is not endorsable by a rational way of making the decision. Taking it to be a rational choice nevertheless, as MAXIMALITY does—*that* takes the cake.

References

Arntzenius, Frank, Adam Elga, and John Hawthorne. 2004. Bayesianism, Infinite Decisions, and Binding. *Mind*, 113(450): 251–283

Bader, Ralf. 2018. Stochastic Dominance and Opaque Sweetening. *Australasian Journal of Philosophy*, 96(3): 498–507

Bales, A., D. Cohen, and Toby Handfield. 2014. Decision Theory of Agents with Incomplete Preferences. *Australasian Journal of Philosophy*, 92(3): 453–470

Bradley, Seamus. 2013. Weak Rationality and Imprecise Choice. *ms*

Doody, Ryan. 2019a. Opaque Sweetening and Transitivity. *Australasian Journal of Philosophy*, 97(3): 559–571

——. 2019b. Parity, Prospects, and Predominance. *Philosophical Studies*, 176(4): 1077–1095

——. 2021. Hard Choices Made Harder. In *Value Incommensurability: Ethics, Risk, and Decision-Making*, edited by H. Andersson and A. Herlitz (Routledge)

Elster, Jon. 1984. *Ulysses and the Sirens* (Cambridge University Press)

Hajek, Alan. 2003. Conditional probability is the very guide of life. In *Probability is the very guide of life: The philosophical uses of chance*, edited by Jr. H. E. Kyburg and M. Thalos (Chicago: Open Court), pp. 187–88

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

Harsanyi, John C. 1955. Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility. *The Journal of Political Economy*, 63(4): 309–321

Hedden, Brian and Daniel Muñoz. 2023. Dimensions of Value. *Noûs*. URL `https://doi.org/10.1111/nous.12454`

Herzberger, H. 1973. Ordinal preference and rational choice. *Econometrica*, 41: 187–237

Levi, Isaac. 1986. *Hard Choices* (Cambridge: Cambridge University Press)

——. 1999. Value Commitmnts, Value Conflict, and the Separability of Belief and Value. *Philosophy of Science*, 66: 509–533

——. 2006. Symposium on Cognitive Rationality: Part I Minimal Rationality. *Mind and Society*, 5: 199–211

——. 2008. Why Rational Agents Should Not Be Liberal Maximizers. *Canadian Journal of Philosophy*, 38(Supplementary Vol 34): 1–17

Schoenfield, Miriam. 2014. Decision Making in the Face of Parity. *Philosophical Perspectives*, 28(1): 263–277

Seidenfeld, Teddy. 1988. Rejoinder [to Hammond and McClennen]. *Economics and Philosophy*, 4(2): 309–315

Sen, Amartya. 1997. Maximization and the Act of Choice. *Econometrica*, 65: 307–312

——. 2004. Incompleteness and Reasoned Choice. *Synthese*, 140(1/2): 43–59

Tversky, Amos and Daniel Kahneman. 1979. Prospect Theory: An Analysis of Decision Under Risk. *Econometrica*, 47: 263–291

——. 1984. Choice, Values, and Frames. *American Psychologist*, 39: 341–350

——. 1986. Rational Choice and the Framing of Decisions. *The Journal of Business*, 59(4 Part 2: The Behavioral Foundations of Economic Theory): S251–S278

Ullmann-Margalit, Edna and Sidney Morgenbesser. 1977. Picking and Choosing. *Social Research*, 44(4): 757–785