PHILOSOPHICAL PERSPECTIVES

Philosophical Perspectives, 28, Ethics, 2014

DECISION MAKING IN THE FACE OF PARITY¹

Miriam Schoenfield The University of Texas at Austin

This paper defends a constraint that any satisfactory decision theory must satisfy. I show how this constraint is violated by all of the decision theories that have been endorsed in the literature that are designed to deal with cases in which opinions or values are represented by a set of functions rather than a single one. Such a decision theory is necessary to account for the existence of what Ruth Chang has called "parity" (as well as for cases in which agents have incomplete preferences or imprecise credences). The problem with the all of the decision theories that have been defended to account for parity is that they are committed to a claim I call unanimity: when all of the functions in the set agree that an agent ought to do A, then an agent ought to do A. A decision theory committed to unanimity violates the constraint I defend in this paper. Thus, if parity exists, a new approach to decision theory is necessary.

1. Introduction

Ruth Chang (1997, 2002, and 2005) has made a strong case for the claim that the familiar trichotomy of value relations that might hold between two comparable outcomes (better than, worse than, equally good) is not exhaustive.² Chang argues that there is a fourth relation, *parity*, which can't be subsumed under any of the others.

The taste of a cup of tea, says Chang (2002), is no better or worse than the taste of a cup of coffee, yet the tea and coffee are not *equally* good. For note that a slight improvement to the tea (say, through the addition a bit of sugar) doesn't break the tie. The taste of tea with a bit of sugar is *still* no better than the taste of coffee and the taste of coffee with a bit of sugar is still no better than the taste of tea. In contrast, four quarters doesn't have more or less value than a dollar bill, but if we increase the value of either, even a teeny bit (say, by adding a penny), the tie will be broken. Four quarters plus a penny *is* more valuable than a dollar bill and a dollar bill plus a penny is more valuable than

four quarters. Consideration of such examples leads Chang to think that there are different value relations that might hold between two comparable outcomes A and B that are consistent with A being no better than B and B being no better than A. They are:

Equality: A and B are of equal value if and only if they are comparable, A is no better than B, B is no better than A, and, for all C, if C is equal in value to one of A or B, it is also equal in value to the other (that is, equality is transitive).

Parity: A and B are on a par if and only if they are comparable, but none of the relations: {better than, worse than, equally good as} hold between A and B.

Note that parity is not transitive since there are cases in which A+ (tea with sugar) is on a par with B (coffee), and B (coffee) is on a par with A (tea), but A+ (tea with sugar) is not on a par with A (tea) – it is better than A. Chang (2005) writes:

The cases of interest are not cases of ignorance, in which one of the traditional three relations holds but we don't know which, nor cases of vagueness, in which the items occupy the borderline of one of the traditional three relations but are cases of determinate comparability; therefore, as cases of determinate comparability in which none of the traditional three relations holds, they must be cases in which a fourth relation of comparability holds — they are on a par. (331)

The existence of parity has been questioned by some,³ but, in this paper, I am going to assume that Chang is right in claiming that parity exists. I am also going to assume, like Chang, that there are attitude-independent facts about value relations, though I suspect that many of the arguments in this paper would apply even if this assumption were dropped.

On the assumption that parity exists, the value facts cannot be represented by a value function that assigns to each outcome a unique real number. This is because trichotomy holds of real numbers: for all real numbers r_1 and r_2 , r_1 is greater than, equal to (where equality is transitive), or less than r_2 . But, if parity exists, trichotomy doesn't hold of evaluative comparisons. Thus, if there are outcomes that are on a par, the value facts cannot be represented by a traditional value function. This is of interest because without a traditional value function we cannot do traditional decision theory.

In traditional decision theory, the *expected* value of an action is the average value of the outcomes that might result from performing the action, weighted by the probability that those outcomes will obtain if the action is performed. Expected value is of particular interest to consequentialist ethics since many consequentialists think that the action one *ought* to perform is the one that maximizes expected value. (At very least, this is true of what is sometimes called the *subjective ought* — the "ought" that takes our ignorance about how the world is into account). But even non-consequentialist ethicists should be interested in expected value since, plausibly, there are *some* cases in which what one ought to do is determined by expected value considerations. Take, for example, a version

of Frank Jackson's (1991) famous pill case. Suppose you are considering which of three pills to give to a child who is extremely ill. You know that one of Pill B or Pill C will cure the child with no side effects, the other will kill him, but you have no idea which pill will do what. You also know that Pill A will cure the child, but a mild headache is a likely side effect. Clearly, you should give the child Pill A, and a plausible explanation for this is that Pill A is the option that maximizes expected value.

The question that now arises is as follows: if we don't have a value function that assigns a number to each outcome, then how is expected value theory to be carried out? An analogous question in epistemology has recently received a great deal of attention. Epistemologists have been concerned with how decision theory should be carried out when one's doxastic state cannot be represented by a single *probability* function, one which assigns to each proposition a number representing one's degree of confidence in that proposition.

As we'll see, very similar strategies are appealed to in the epistemic and evaluative cases to represent the probabilities and values in non-standard ways. If it's plausible that it can be rational to hold opinions that fail to be representable by a single probability function⁴ and it's plausible that the value facts cannot be represented by a single value function, then it's plausible that sometimes we'll face situations where *neither* the relevant probabilities *nor* the relevant values are representable with the kinds of functions that feature in traditional decision theory. Thus, the project of extending decision theory to deal with cases in which *opinions* can't be represented precisely and the project of extending decision theory to deal with cases in which *values* can't be represented precisely should be pursued in tandem.

I will not, in this paper, provide a decision theory to account for parity. Rather, I will defend a *constraint* that such an extended decision theory must satisfy and I will show that what has become the standard approach in both the epistemic and evaluative cases fails to satisfy this constraint. My conclusion will be that if we want to do decision theory without traditional value and probability functions, a different approach is necessary.

2. Modeling Parity

The first step in developing a decision theory that accounts for parity is to come up with an alternative formal representation of the value facts. Gert (2004) and Chang (2005) describe a variety of ways to use *intervals* (rather than real numbers) to represent parity. A suggestion along the lines of Aumann (1962) would use *vectors*, and one along the lines of Hare (2010) would use a *set* of value functions rather than a single one. Intervals, vectors, or sets of functions can be helpful because they allow for a richer array of ways in which values can be related. Let's look at the set of functions model. On this model the following biconditionals hold:

266 / Miriam Schoenfield

- (a) A is better than B if and only if all of the value functions in the set assign higher value to A than to B.
- (b) A is worse than B if and only if all of the value functions in the set assign lower value to A than to B.
- (c) A and B are equal in value if and only if all of the value functions assign equal value to A and B.
- (d) A and B are on a par if and only if the value functions in the set disagree over the relative values of A and B.

This model allows for a case in which A+ is better than A (all of the functions assign higher value to A+ than to A), A and B are on a par (the functions disagree about the relative values of A and B), and A+ and B are also on a par (the functions disagree about the relative values of A+ and B). A very similar strategy can be carried out using vectors⁵ or intervals⁶ but, for ease of discussion, I will focus exclusively on sets of functions. None of the main points that I argue for will rely crucially on the use of sets, rather than vectors or intervals.

Once we have a model for representing values, we need to think about how this model will play into a decision theory. There have been a variety of rules proposed about how decision theory should work when we use sets of value functions, rather than a single function, but all of the views endorsed in the literature⁷ are committed to:

Value-Unanimity: (When value considerations are the only relevant considerations then) if an agent is faced with a choice between A and B, and, according to each of the value functions in the set of functions representing the value facts, the expected value of A is greater than the expected value of B, she (subjectively) ought to do A.8

A similar approach has also been applied to represent what epistemologists call "imprecise credences." Rather than representing an agent's opinions with a single credence function we can use a set of credence functions (called a "representor"). The literature on decision theory using imprecise credences contains a variety of different decision rules. But all of the rules that have been proposed 10 are committed to:

CREDAL-UNANIMITY: (When value considerations are the only relevant considerations then) if an agent is faced with a choice between A and B, and, according to every credence function in the agent's representor, the expected value of A is greater than the expected value of B, she (subjectively) ought to do A.

Note that Value-Unanimity assumes a single credence function and Credal-Unanimity assumes a single value function. But, as I mentioned earlier, if there is parity, and a rational agent's opinions aren't representable by a single credence function, our decision theory should account for cases where *neither* the values nor the probabilities are representable by a single function. The weakest generalization of Credal or Value Unanimity, which accounts for multiple

functions of each kind, says that if *every* possible combination of credence and value functions requires doing A, then an agent ought to do A. This gives us:

UNANIMITY: Let S be a rational agent facing two options: A and B. S's opinions are represented by a set of credence functions, C, and the value facts are represented by a set of value functions, V. If, according to every $\{c,v\}$ pair, where $c \in C$ and $v \in V$, the expected value of A is greater than the expected value of B, and value considerations are all that's at stake, an agent (subjectively) ought to do A.

The aim of this paper is to argue against UNANIMITY.

3. The First Premise: LINK

My argument against UNANIMITY is based on a principle I will call "LINK." 11

LINK: In cases in which considerations of value are the only ones that are relevant, if you are rationally certain that one option, A, will bring about greater value than the alternative option, B, you're required to choose A. If you are rationally certain that neither of the two options will bring about greater value than the other, it's not required that you choose A, and it's not required that you choose B.¹²

If we want to extend the traditional notion of expected value to govern the subjective ought in cases of parity, such that what an agent subjectively ought to do is what maximizes expected value (when value is all that's at stake), then LINK entails the following constraint on the notion of expected value:

Link_{EV}: Let A and B be two options and let p be an agent's rational probability function. For any outcomes X and Y, let "V(X)>V(Y)" stand for the proposition that the outcome that would result from choosing X is better than the outcome that would result from choosing Y. Then,

- (1) If p(V(A) > V(B)) = 1, then $EV^p(A) > EV^p(B)$,
- (2) If p(V(A) > V(B)) = 0 and p(V(B) > V(A)) = 0, then $\sim (EV^p(A) > EV^p(B))$ and $\sim (EV^p(B) > EV^p(A))$

Since Link and Linkev are so closely connected I will sometimes refer to Linkev as just "Link."

LINK is so named because it is a principle that puts constraints on *expected value* (or the subjective *ought*) given the knowledge an agent has about *value* (which, for a consequentialist, tracks the objective *ought*). Traditional, single-function expected value theory satisfies LINK and I'm going to argue that *any* theory of expected value should satisfy LINK. I will then argue that decision theories committed to UNANIMITY must reject it. Hence, decision theories that are committed to UNANIMITY should be rejected.

The main argument of this paper can thus be summarized as follows:

- (1) LINK is true.
- (2) LINK is inconsistent with UNANIMITY.
- (3) UNANIMITY is false.

Why accept Link? Because, if Link is rejected, expected value theory cannot play the role that it was intended to play: namely, providing agents with limited information guidance concerning how to make choices in circumstances in which value-based considerations are all that matter. It is important to realize that the *ultimate* good, for consequentialists, or in cases in which consequentialist reasoning is appropriate, is an outcome that is valuable, not an outcome that is expectedly valuable. A consequentialist is only interested in *expected* value insofar as it helps her obtain what is actually important: value. Given this, value and expected value cannot float free of one another. Any theory of expected value that makes demands that don't make sense given our concern with *value* can't do what expected value theory is meant to do.

To see why rejecting LINK undermines the role that expected value theory is meant to play, suppose first that there are cases in which an agent is rationally certain that V(A) > V(B) and yet EV(B) > EV(A). If such a case exists, it is one in which the agent's concern with *value* warrants choosing A (she knows that A will bring about a better outcome!), but if she is to maximize *expected* value she must choose B. Thus, in such a case, the demand to maximize expected value yields a recommendation that doesn't make sense given an agent's concern with *value*. If an agent knows that V(A) > V(B) but EV(A) is neither greater nor less than EV(B) then an agent's concern with *value* warrants choosing A. But if she is aiming to maximize *expected value* she will be indifferent between A and B.

It had also better not turn out that an agent can be in a situation in which she knows that neither V(A) nor V(B) is greater than the other, yet the expected value of one is greater than the other. If expected value theory required us to make choices that we are certain would lead to no improvement in value, then expected value theory is imposing requirements that transcend what we actually care about: the achievement of *value*. (If it's known that eating kale won't make you healthier than eating spinach and eating spinach won't make you healthier than eating kale, then, if your doctor recommends spinach over kale, she is imposing constraints that transcend her role of providing good guidance given your concern with *health*).

Here is another way of making the point: expected value theory is designed to help us in our pursuit of value under conditions of *uncertainty*. Frequently, under conditions of uncertainty, we don't know which of our actions will produce the outcome with the highest value. But in cases in which there is *certainty* about the relative values of the available options, what one ought to do to maximize *expected* value should line up with one ought to do to maximize *value*.

In sum, if you're developing a theory of expected value which will govern your decision making under conditions of uncertainty, and your interest in developing this theory stems from the importance of producing outcomes that have high *value* (and avoiding those that have low *value*), then you need to make sure that your theory of *expected* value respects your interest in *value*. A theory which violates LINK is a theory with an expected-value fetish. It mistakenly requires or permits actions that shouldn't be required or permitted given an agent's concern with value.

4. The 2nd Premise: LINK conflicts with UNANIMITY

In this section, I appeal to *opaque sweetening cases*¹³ to show why LINK conflicts with UNANIMITY. To begin, consider the following example:

CHARITY TRANSPARENT

You have come into a large sum of money and you want to donate it to charity. You have two options. Either you could expand the free meal program for the homeless to include an extra 50 meals a day or you could start an arts program in a low-income school.

Let us suppose that you have no associative obligations or personal preferences about which organization to help, and that your personal values, commitments, and considerations of justice don't favor one or the other. Additionally, neither option involves performing an action that is wrong in itself (breaking a promise, violating someone's rights, and so on). Let us assume, then, that the *only* relevant considerations are value based¹⁴ and that you know that the outcomes of donating to each charity are on a par.

Now consider the following version of the charity case:

CHARITY OPAQUE

As before, you have a choice between donating to the meals program or the arts program. But this time you are in a rural area and you need to deliver the money in cash to the office of the program that you decide to donate to. You know that one is east and one is west, but you don't remember which is which. You assign a 0.5 credence to the meal program being east and a 0.5 credence to the arts program being east. You are in the middle of nowhere, with no phone or internet access, and you only have a couple of hours before you need to start traveling to the airport. "Well," you think to yourself, "so much for all these complicated considerations. I guess I'll just choose which direction to go randomly." As you are getting into your car, still undecided about where you are going, you see in your rear view mirror a man with a suitcase, heading eastward, most likely to the train station. You suspect that the traveler could make it to the train station just fine on his own, but it is a hot day and the case looks heavy. You think to yourself: "Hmm... if I went east I could help the traveler out by giving him a ride."

Question: what should you do? As I'll argue, LINK entails that it's permissible for you to go either east or west, but UNANIMITY requires that you go east. Thus LINK and UNANIMITY are incompatible.

Let's begin by applying UNANIMITY. Although the set of functions representing the facts about value might include some functions which assign a higher value to the meal program than to the arts program and others which assign a higher value to the arts program than to the meal program, all of the functions will agree that going east has a greater expected value than going west. Why is this?

Take any function $v \in V$ and let us suppose that v assigns the following values:

```
O_1 : You donate to the meal program and don't give the traveler a ride. \textit{v}(O_1)
```

= x

 O_2 : You donate to the arts program and don't give the traveler a ride. $\nu(O_2)$ = y

Since the value of donating to the meal program and helping the traveler is greater than the value of just donating to the meal program, the set of functions model requires that every value function assign greater value to the meals program + helping the traveler than to just the meals program. Similarly for the arts program. Thus, we have:

```
O_{1+}: You donate to the meal program and give the traveler a ride. \nu(O_{1+}) = x+r where r>0.
```

 O_{2+} : You donate to the arts program and give the traveler a ride. $\nu(O_{2+})$ = y+s where s>0.

Recall that we are assuming that your credence that the meal program is east equals your credence that the arts program is east = 0.5. So if we use v to calculate expected values we have:

```
EV(Going east) = (.5)(x+r) + (.5)(y+s).
EV(Going west) = (.5)(x) + (.5)(y).
```

Since r and s are positive, the expected value of going east is greater than the expected value of going west. Since v was an arbitrary function amongst the set of functions representing the value facts, we can conclude that it is true of *all* functions in V that the expected value of going east is greater than the expected value of going west. Unanimity thus entails that you ought to go east.

But Link entails that it is permissible to go east *or* west. This is because you know that the value of the outcome that would result from going east is no greater or less than the value of the outcome that would result from going west. After all, you know that the meal program and the arts program are on a par. Thus, regardless of which program is east, adding to that program the value of

assisting the traveler doesn't break the tie. Since you know that the value of going east isn't greater than the value of going west, and the value of going west isn't greater than the value of going east, LINK says that either option is permissible.

Note that all that is necessary to generate the conflict between LINK and UNANIMITY is that there exist cases in which:

(a) you are faced with a choice between two outcomes, O_1 and O_2 , which are on a par

and

(b) value considerations are the only relevant considerations.

We can modify any such case so that (1) you are now faced with two choices, A and B, (2) you know that choosing A will bring about O_1 or O_2 and choosing B will bring about O_1 or O_2 , but you don't know which choice will bring about which outcome, and (3) you know that a little bit of extra value will be added if you choose A.¹⁶

This completes the argument against Unanimity. Premise (2), which says that Link and Unanimity are incompatible, seems undeniable. But perhaps Premise (1) (Link) should be rejected. I describe some motivations for rejecting Link and respond to these in the next section.

5. Rejecting LINK

It may seem intuitive that you should go east and help the traveler in the opaque case. If so, you might think that this judgment motivates a rejection of LINK, since LINK entails that going west is permissible. Hare (2010) is sympathetic to the thought that, in cases like these, you should help the traveler and he motivates the judgment with the following principle:

MOST-REASON: If all of your (subjective) reasons¹⁷ to do B are also reasons to do A, but you have an *additional* (subjective) reason to do A that isn't a reason to do B, you have most reason to do A.

Most-Reason certainly has some intuitive appeal. It's hard to see how you could fail to have most reason to do A when you have a reason to do A that isn't a reason to do B, but *all* of your reasons to do B are also reasons to do A. Let's apply the principle to the opaque charity case. It certainly looks like, in the opaque case, every (subjective) reason you have to go west is also a (subjective) reason to go east. Since you don't know which program is in which location, your only reason to go west is the fact that, by going west, you will help *either* the meals program or the arts program, and this is also a reason to go east. Furthermore, you have an *additional* reason to go east which is not a reason to go west, namely, that by doing so you will be able to help the traveler. Since all

of your reasons to go west are reasons to go east, but you have an additional reason to go east, which is not a reason to go west, it follows that you have *most* reason to go east.

Note that this argument could not be given if you knew which charity was at which location. For then it would not be the case that every (subjective) reason you have to go west is a reason to go east. If you knew, for instance, that the arts program was west, then you would know that by going west you could give students an opportunity for artistic expression. This would be a subjective reason to go west that is not a subjective reason to go east.

At this point you might think that, rather than simply assuming that LINK is true and that UNANIMITY should be rejected, we should recognize that there are important considerations on both sides. Thus, rather than an argument, we have a puzzle. On the one hand, LINK seemed like a plausible principle about how value and expected value interact. But there are at least two points of pressure in the other direction: (1) Our intuitions about Charity-Opaque may favor a requirement to help the traveler and (2) Most-Reason, itself a plausible principle, favors helping the traveler. Some philosophers might think that if we add to these considerations (3) the fact that Unanimity allows us to construct a nice decision theory in the face of parity, and one that sits well with the work that has been done in epistemology to deal with similar issues, the balance of reasons favors rejecting Link.

Let me respond to each of these concerns in turn, beginning with the intuitive motivations. The first thing I'd like to note is that people have differing intuitions about the case. Second, even if the intuition is sufficiently robust, there are plausible ways of explaining it away. For example, there are a variety of dimensions of moral evaluation that could, in principle, be appealed to in explaining why there is *something* non-morally ideal about going west, even if doing so is permissible. Perhaps going west reveals bad character. Or perhaps going east is an action with greater moral worth. Of course, more would have to be said to make these hypotheses plausible but, for now, I just want to point out that there are other avenues of criticism that are consistent with the judgment that going west is *permissible*. Additionally, the intuition is plausibly driven by the fact that usually when we know that one of two outcomes will come about by choosing A, and another will come about by choosing B, but we don't know which is which, if an improvement is added to A, we should choose A. For as long as we leave open the possibility that A and B are equally valuable or that one is more valuable than the other, a plausible decision theory that is consistent with LINK would recommend choosing what Hare has called the "sweetened" option. It is only in cases in which we know that the two outcomes are on a par that there's no reason to take the sugar. 18 Given that people differ about their intuitions in opaque cases, and that there are ways of explaining away these intuitions, such intuitions should not be thought as decisive. This is certainly true given that, as I've argued, decision theory cannot play its intended role if LINK is rejected.

The second consideration offered on behalf of helping the traveler was Hare's argument appealing to Most-Reason. But this argument is problematic. First, for the argument to apply we have to assume that all of the your reasons for going west are also reasons for going east. But whether this is so depends on how reasons are individuated. If you think that your only reason for going west is that by doing so you could aid the meal program or the arts program, that is, indeed, a reason to go east. But if you think that your reason for going west is that you could help whichever program is located in the west, that is not a reason to go east. 19 Second, Most-Reason is a questionable principle. Reasons interact in complex ways and they don't always add up as one might expect them to. For example, Jonathan Dancy (2004) talks about cases in which you gain a reason to do some action, A, which attenuates the strengths of some of your other reasons to do A. So consider a case in which, initially, all of your reasons to do B are reasons to do A. You then gain an additional reason to do A, but that reason attenuates the strength of the initial reasons you had to do A. In such cases, you won't always end up with most reason to do A. Indeed, it shouldn't be surprising that cases in which parity is involved will be cases in which additional reasons don't always translate into most reason. After all, one of the distinctive features of outcomes that are on a par is that small improvements don't break the tie. Finally, Bales, Cohen and Handfield (2014) provide some nice examples, which I won't rehearse here, that demonstrate further problems with Most-Reason.

The last concern I want to address involves weighing the plausibility of LINK against the desirability of the nice decision theories that have been constructed around Unanimity. I think that rejecting LINK on the grounds that it is inconsistent with some otherwise attractive decision theories would be misguided. For without LINK, it becomes unclear why we are *interested* in decision theory to begin with. Why should be concerned with maximizing some quantity: "expected value" in cases in which we know that doing so will yield absolutely no benefits in terms of *what we actually care about*: value. It seems to me that no matter how elegant a decision theory you can construct with Unanimity, since it requires rejecting LINK, what we're left with is an elegant theory that doesn't track anything of actual importance.

6. Conclusion

In this paper I argued that a satisfactory decision theory must reject UNANIMITY. This is because UNANIMITY is inconsistent with LINK, and no decision theory that rejects LINK can play the role that decision theory is designed to play: guiding agents in their pursuit of value under conditions of uncertainty. Since UNANIMITY is assumed in all of the decision theories that account for parity that have been endorsed in the literature, if parity exists, a new kind of decision theory is necessary.

274 / Miriam Schoenfield

The rejection of UNANIMITY also has implications for epistemology. All of the decision theories that have been proposed for agents with imprecise credences are committed to CREDAL-UNANIMITY. Even those decision theories that are specifically designed to *avoid* problems with decision theories for imprecise agents, like the problems described by Elga (2010), are committed to CREDAL-UNANIMITY. CREDAL-UNANIMITY, if generalized to account for multiple value functions, yields UNANIMITY. Since UNANIMITY must be rejected, theorists interested in decision theories for agents with imprecise credences must also take a different strategy.

The only proposal I am aware of that is *not* committed to UNANIMITY is one that is described, though ultimately rejected, by Hare. The decision theory that Hare describes involves some additional machinery and evaluating its merits is beyond the scope of this paper. Bales et al. think that there are some problems with the proposal²⁰ but, given that it is the only decision theory that has been described that is consistent with Link, it is probably a good place to start.

I'd like to end by raising the possibility that the difficulties surrounding parity should simply lead us to reject its existence. There are two ways to do this. First, one might think of Chang's cases as cases where there is *indeterminacy* about which of two outcomes is better, rather than cases in which *determinately* some fourth relation (parity) holds between them.²¹ While there may be good reasons to take this route, it's not clear to me that this strategy would help with the problem I describe in this paper. For even if we substitute indeterminacy for parity, UNANINIMITY will say that in the opaque charity case it is *determinately* true that one ought to go east even though one is *certain* that neither option is determinately better than the other. This seems problematic for the sorts of reasons I appealed to do in defending LINK. If you know that it's indeterminate what the relative value of the two options before you is, and what you care about is *value*, how can it be *determinately* true that one is obligatory?

The second way to reject the existence of parity is to suppose that the facts about value *can* be represented by a precise value function. Like Chang, it strikes me as implausible that the facts about value provide us with the kind of precise tradeoffs that would be necessary. *Exactly* how many days of pleasant weather is a mediocre piece of art worth? How many bunnies' lives would it be worth sacrificing to forever eliminate the common cold? I don't think that the evaluative nature of the world provides us with answers to these questions. But if it did, decision theory would certainly be a lot easier.

Notes

 I would like to thank Jonathan Dancy, Sinan Dogramaci, Tom Dougherty, Caspar Hare, Elizabeth Harman, Luke Elson, Julia Markovits and especially Susanna Rinard for incredibly helpful conversations on this topic and for reading and commenting on earlier drafts. I am also indebted to Adam Elga and the members

- of his Spring 2014 Decision Theory seminar for a very helpful discussion of the paper.
- By comparable I just mean that value comparisons between the two outcomes make sense. This rules out cases of complete incommensurability like those described in Raz (1985).
- 3. See, for example, Gert (2004) and Gustaffson and Espinoza (2010). For a discussion of an alternative explanation of the small improvement cases in terms of vagueness see e.g. Broome (1997), and Elson (ms.). For an epistemic analogue of this view see Rinard (forthcoming).
- 4. This claim has challenged by some. See for example Elga (2010) and White (2010).
- 5. The vector space gives an incomplete ordering of outcomes, and sets of value functions can be defined over the vector space by generating functions from the vector space to the real number line, each corresponding to a different completion of the ordering. Thus, we can use the vector space to derive a set of value functions, and cases in which the value functions disagree about the relative values of A and B, will be the instances of incompleteness in the initial ordering.
- 6. See Chang (2005) for details.
- 7. See for example Chang (2005), Hare (2010), Weirich (2004, p.67) and Good (1952, p.114). Hare does describe a view that is not committed to (and is, indeed, inconsistent with) VALUE-UNANIMITY but, ultimately, he finds the view he calls "Prospectism" (which *is* committed to VALUE-UNAMINITY) preferable.
- 8. I am assuming that the agent's credences are rational and that the agent is not ignorant of the value facts. Dealing with normative uncertainty brings with it a whole host of interesting issues that I will not get into here. For a discussion of such issues see, for example, Sepielli (2013).
- 9. See, for example, Levi (1974) and Joyce (2010).
- 10. See, for example, Elga (2010), Joyce (2010), Sud (2014) and Rinard (forthcoming).
- 11. A different constraint on decision theory that turns out to be incompatible with UNANIMITY is argued for by Bales, Cohen and Handfield (2014). Bales et al. also use opaque sweetening cases (which will be discussed later, and were first described in Hare (2010)) to show how their constraint is violated by the decision theory that they are targeting: Hare's prospectism. My argument and theirs should be seen as complementary and there are some important differences between them that are worth noting. First, I defend the main premise of my argument, LINK, by appealing to the role we want expected value theory to play. The defense of their central principle, Competitiveness, is based on its similarity to a dominanceprinciple and an appeal to case judgments (see p.7–8). (The kind of dominance principle that they appeal to faces concerns from Nozick (1969). Bales et al. are aware of these worries and add some stipulations to avoid them). One reason that it is desirable to have an argument against UNANIMITY that is not based on intuitive judgments (about cases, or principles) is that, as we'll see, there will be intuitions favoring both the acceptance and the rejection of the principle. Rather than trying to adjudicate between these intuitions, my approach appeals to a different kind of consideration altogether: the role that decision theory is meant to play. This consideration, I will argue, is overriding. Additionally, part of my argument will involve showing how the sorts of considerations favoring LINK have important analogues and implications in epistemology.

276 / Miriam Schoenfield

- 12. Two notes about the formulation of the principle: First, if you prefer to think of rational decision making as determined fully by the agent's preferences rather than objective facts about value, simply substitute "A will bring about greater value than the alternative B" with "A will bring about an outcome that is preferred to B." Second, the principle can be stated in terms of rational certainty or knowledge. I will just assume that, in the cases we are considering, certainty is rational if and only if the agent knows and so I won't differentiate between knowledge and rational certainty. This is a harmless assumption since even if we restrict ourselves to cases in which the agent *both* knows and is rationally certain, a restricted version of the principle that applies only to such cases will still cause problems for UNANIMITY.
- 13. Opaque sweetening cases in the practical realm are first described in Hare (2010) and have also been discussed by Bales et al. (2014). Schoenfield (2012) discusses epistemic opaque sweetening cases.
- 14. If you think that this *isn't* the kind of case in which what one subjectively ought to do is determined by value-based considerations, feel free to substitute your own. (For example if you think that interpersonal comparisons of value don't make sense, you can imagine a case where you're considering two different ways of helping the same group of people).
- 15. Fill in the case however you like to make a 0.5 credence reasonable. Perhaps the locations were determined by the flip of a fair coin.
- 16. I am assuming that, in at least some cases satisfying (a) and (b), if value is all that is at stake in the transparent version of the case, this remains true when you no longer know which outcome will result from your choice.
- 17. See Vogelstein (2012) for a discussion of subjective reasons. For our purposes it will suffice to think of subjective reasons as reasons that the agent is aware of.
- 18. I borrow this phrase from Hare as well.
- 19. Thanks to an anonymous referee and to Adam Elga for discussion on this point.
- 20. See their note 12.
- See Broome (1997) and Elson (ms.). Rinard (forthcoming) develops a similar view for imprecise credences.

References

- Aumann, R. (1962). "Utility Theory without the Completeness Axiom." *Econometrica* 30(3): 445–462.
- Bales, A., Cohen, D. and Handfield, T. (2014). Australasian Journal of Philosophy 92(3): 453-470.
- Broome, J. (1997). "Is Incommensurability Vagueness?" In Chang, R. (ed.) *Incommensurability, Incomparability and Practical Reason*. Harvard University Press.
- Chang, R. (1997) Introduction to Incommensurability, Incomparability and Practical Reason Harvard University Press.
- Chang, R. (2002). "The Possibility of Parity." Ethics 112(4): 659–88.
- Chang, R. (2005). "Parity, Interval Value, and Choice." Ethics 115(2): 331-50.
- Dancy, J. (2004). Ethics Without Principles. Oxford University Press.
- Elga, A. (2010). "Subjective Probabilities should be Sharp." Philosopher's Imprint 10(5).
- Elson, L. (ms.) "Why not Incommensurability as Vagueness?"

- Gert, J. (2004). "Value and Parity." Ethics 114(3): 492-510.
- Good, I.J. (1952). "Rational Decisions." Journal of the Royal Statistical Society 14(1): 107-14.
- Gustafsson, J. and Espinoza, N. (2010). "Conflicting Reasons in the Small-Improvement Argument." *The Philosophical Quarterly* 60 (241): 754–763
- Hare, C. (2010) "Take the Sugar." Analysis 70(2): 237-247.
- Jackson, F. (1991) "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection." Ethics 101(3):461–482.
- Joyce, J. (2010). "A Defense of Imprecise Credences in Inference and Decision Making." Philosophical Perspectives 24(1): 281–323.
- Levi, I. (1974). "On Indeterminate Probabilities." Journal of Philosophy 71(13): 391-418.
- Nozick, R. (1969). "Newcomb's Problem and Two Principles of Choice." In Rescher, N. (ed.). *Essays in Honor of Carl Hempel.* D. Reidel.
- Raz, J. (1985). "Value Incommensurability: Some Preliminaries." Proceedings of the Aristotelian Society 86: 117–134.
- Rinard, S. (forthcoming) "A Decision Theory for Imprecise Credences." Philosophers Imprint.
- Schoenfield, M. (2012) "Chilling Out on Epistemic Rationality: A Defense of Imprecise Credences." *Philosophical Studies*. 158(2): 197–219.
- Sepielli, A. (2013). "What to do when you don't know what to do..." Noûs 47(1): 521-544.
- Sud, (2014). "A Forward Looking Decision Rule for Imprecise Credences." Philosophical Studies 167(1):119–139
- Vogelstein, E. (2012). "Subjective Reasons" Ethical Theory and Moral Practice 15(2):239-57.
- Weirich, P. (2004). Realistic Decision Theory: Rules for Nonideal Agents in Nonideal Circumstances. Oxford University Press.
- White, R. (2010). "Evidential Symmetry and Mushy Credence." Oxford Studies in Epistemology Vol. 3. Tamar Szabo Gendler and John Hawthorne eds. Oxford University Press.