

## ENDS TO MEANS

### A PROBABILITY-RAISING ACCOUNT OF MEANS AND THE WEIGHT OF REASONS TO TAKE THEM

Matthew S. Bedke

LET US SAY THAT YOUR ENDS are *whatever* you have ultimate (underived) reason to do or to bring about. This leaves open whether your ends are stance dependent and so “given” to you by your contingent desires, your nature as a rational being, your self-identity, etc., or whether at least some of your ends are not stance dependent and so not “given” to you in any of these ways.

So understood, it is uncontroversial that you have reason to take the means to your ends. More specifically, some reason to do or bring about an end is going to transmit to reason to take the means. Using this as a point of departure, this paper considers *what it is* to be a means to an end and *how much* reason transmits from an end to its means. The theory on offer is a probability-raising theory that says roughly this: an action is a means to an end just in case it raises the probability of the end relative to *the worst one could do*—i.e., relative to that action that would make the end least probable. And the amount of reason transmitted from an end to a given means is a function of the degree to which it raises the probability of the end relative to the worst one could do.

Some recent literature addresses these issues, though given the importance of means-ends relations and the ascendance of reason-based analyses of normativity the issues remain relatively underexplored.<sup>1</sup> My own exploration will proceed

- 1 The account here is meant to improve upon and replace the Transmission Principle in Bedke, “The Iffiest Oughts.” For similar probability-raising accounts, see Kolodny, “Instrumental Reasons”; and Stegenga, “Probabilizing the End.” For a very different approach, see Horta, “Reasons as Defaults.” For the related issue of what it is to promote the satisfaction of one’s desires, see Lin, “Simple Probabilistic Promotion”; Sharadin, “Checking the Neighborhood” and “Problems for Pure Probabilism about Promotion (and a Disjunctive Alternative)”; Coates, “An Actual-Sequence Theory of Promotion”; Schroeder, *Slaves of the Passions*; and Finlay, *Confusion of Tongues*. For a skeptical take on the probability-raising approach to promotion, see Behrends and DiPaolo, “Finlay and Schroeder on Promoting a Desire” and “Probabilistic Promotion Revisited.” The present project is related to those discussions of means-ends reasons and reasoning that focus on necessary means (see, e.g., Bratman, “Intention, Practical Rationality, and Self-Governance,” 424; Darwall, *Impartial*

by generating plausible proposals within a certain framework and “Chisholming” to a correct analysis. Though I do not offer a comprehensive compare and contrast with other approaches in the literature, along the way I indicate some of my reservations about some of those approaches.

# 1. WHY PURE PROBABILITY RAISING?

Before I begin, let me briefly address some objections to a pure probability-raising approach, which aspires to offer necessary and sufficient conditions for being a means and to provide resources for modeling reason transmission. I hope I can do this in a helpful way before nailing down the best version of a probability-raising theory and answering some intramural questions within that approach.

One big-picture issue is whether probability raising of any sort is necessary for being a means. The crucial question here is whether there are means to an end that do not raise the probability of the end. In that vein, consider the following. Ann is golfing and one of her ends is to get a hole in one. If she slices the ball, intuitively that reduces the probability she will get a hole in one relative to some relevant baseline (e.g., hitting a clean shot).<sup>2</sup> But on this occasion, Ann gets lucky. Though her tee shot slices the ball, she hits it *just so* and the result is a hole in one. Arguably, this is a case in which an action (slicing the ball) is a means to an end (getting a hole in one) by virtue of causing it (/bringing it about) and despite the fact that it lowered the probability of that outcome relative to a relevant baseline.<sup>3</sup> If so, probability raising is not necessary, and being an action that causes, brings about or constitutes an end is sufficient for being a means.

Cases like this raise tricky issues. One question is how to demarcate an action so that we might ask of *it* whether it is a cause or a probability raiser or a means to an end. For instance, are we to consider whether *slicing the ball* is a means to the hole in one, or *slicing the ball just so* is a means to the hole in one? It looks like the coarse-grained action of slicing (any which way, as it were) has less claim to being a cause of the end than the more fine-grained action of slicing it just so, for there are lots of ways in which the coarse-grained action can play out and very

---

*Reason*, 16; Schroeder, “Means-End Coherence, Stringency, and Subjective Reasons,” 245) or sufficient means (see, e.g., Raz, “Simple Probabilistic Promotion,” 9, and “The Myth of Instrumental Rationality,” 5–6), but the scope here is intended to be broader.

2 Articulatng the relevant baseline is one of the intramural issues that needs to be sorted. But I think the case raises a nice structural question that we can address without first settling on a baseline.

3 The example is adapted from discussions of causation as probability raising in Hitchcock, “Do All and Only Causes Raise the Probabilities of Effects?” and Suppes, *A Probabilistic Theory of Causality*, 41.

few of them result in a hole in one, whereas the fine-grained action leaves little room for ways of doing it that do not result in a hole in one. So let us first consider a specification of the case that focuses on the fine-grained action.

Is the fine-grained action a probability raiser? Arguably, yes. Indeed, there are plausible, objective theories of probability that will assign conditional probability one to an effect given its cause.<sup>4</sup> More generally, there are ways of theorizing about probability raising that can be applied to the golfer's case, which presents an interesting theoretical choice: either plug in a theory of probability that makes slicing the ball *just so* (and other acts that cause/bring about/constitute some end) a probability raiser and thereby defuse the counterexample, or choose a theory of probability under which the fine-grained action fails to be a probability raiser and thereby give the counterexample some teeth. Faced with this choice, the probability-raising approach looks like an attractive, simple, theoretical option flexible enough to handle the case, and not necessarily an approach that succumbs to a clear, devastating counterexample.

On the other hand, it is not clear to me that we should be focusing on the fine-grained action to begin with. For our purposes, a means needs to be an action over which we exercise some *agency* and over which we have some *control*. If Ann is like most of us, she does not have the requisite control over how she slices a golf ball to be able to slice it *just so*. The fine-grained specification of the case does not feature an item eligible for being a means to an end and once again we defuse the counterexample. Note that the problem here is not one of outcome luck. The problem is whether Ann has the agentive ability to produce the requisite intentions and motor functions—the ones that guide slicing the ball *just so*, resulting in a hole in one—rather than nearby intentions and motor functions.

So let us consider a specification of the case that focuses on the coarse-grained action. Though slicing the ball (any which way, as it were) plausibly lowers the probability of the end relative to some relevant baseline, we have already noted that it has a much weaker claim to causing the end. Insofar as we undermine the case for causation, we also undermine the intuitive sense that this action is really a means to the end. Once again, we do not have a clear counterexample—we do not have a clear case of a means that is not a probability raiser. At the very least, these considerations make it difficult to craft a case of an action i) over which

4 I will not have much more to say about which theory of probability we should use, for it could be that different theories of probability serve different contexts. It could be, for example, that some theory of subjective probability is best for a first-person deliberative context, or for assigning blame or praise, or grading character, while some objective theory is best for giving advice or evaluative assessment without blame. Having said that, I have written this paper with some notion of objective probability in mind.

one has agentic control (ii) that causes, brings about or constitutes an end while (iii) resisting a probability-raising analysis.

Suppose we can craft the relevant counterexample that calls into question the necessity of probability raising. Where would that leave us? I think it would leave us with a disjunctive theory of means and their weights. We would say that one set of conditions is sufficient to be a means—causing, bringing about or constituting an end—but we would need a different set of conditions for those means that do *not* cause or bring about or constitute an end. We would also need some account of how reasons transmit in the two cases, and it too might be disjunctive. By contrast, a probability-raising approach promises a unified account of what it is to be a means (probability raising) that dovetails with a unified account of weight transmission (probability raising). So even if we could generate a nice counterexample, we would need to decide between the strength of the counterexample and the relative simplicity and elegance of the probability-raising approach.<sup>5</sup> While adjudication of that issue must await the convincing counterexamples to probability raising and development of the best probability-raising theory (something I try to do below), I hope I have said enough to justify an approach that takes probability raising to be a necessary condition on being a means.

The other big picture issue is whether probability raising is sufficient for being a means to an end. Kolodny offers a case in which a boxer's end is to land a punch and the boxer has a "tell" each and every time he is about to throw a punch.<sup>6</sup> Intuitively, a tell like this seems to be more of an *indicator* or a *pre-signal* that a punch is coming and less like a means to landing a punch *even if* the prob-

5 Lin ("Simple Probabilistic Promotion") has a different argument against causation as a sufficient condition. His argument focuses on cases in which one can cause an end by doing an action, like pressing a button, but the end will occur even if one does nothing (e.g., someone else ensures that the end is brought about even if one does nothing). It seems that *any* amount of reason to do nothing (e.g., pressing the button will produce a mild electric shock) will outweigh the reason to press the button, for the end will occur regardless. This, Lin argues, shows that one does not have *any* reason to press the button, otherwise it would have some weight capable of outweighing some small reason to do nothing. I am not convinced by the argument. One who thinks causing an end (/bringing it about/constituting it) is sufficient to be a means could say that, in these cases, the end transmits *just as much reason to do nothing as it transmits reason to press the button*. After all, the end occurs regardless of what one does. If so, we predict that any *other* reason to do nothing (e.g., the electric shock) will *combine* with the end-given reason to do nothing, and these together will outweigh the end-given reason to press the button. That said, like Lin I reject causing an end as sufficient for being a means, but based on the parsimony considerations in the main text.

6 Kolodny, "Instrumental Reasons."

ability of landing a punch given the tell is greater when compared to a relevant baseline.

Once more, cases like this raise tricky issues. First, we should reiterate a point made above: items eligible to be a means are actions that involve agency and control. So we should rule out interpretations of the boxer's case in which the tell lacks agentive control. Having said that, there could be interpretations of this case or cases like it in which the tell is subject to agentive control, but where it still seems to fall short as a means. So we should further clarify, second, that an action is not a probability raiser for an end simply because the conditional probability of the end given the action is higher than some relevant baseline. To be a probability raiser the action must *make* the probability of the end higher than some relevant baseline—the probability of the end has to be higher at least partly *in virtue of* the action and its relations to the relevant baseline and the end.<sup>7</sup> At least, this is a feature of the present approach. Even if the boxer's tell is subject to agentive control, intuitively it is not helping to *make* the probability of landing a punch higher (the throwing of the punch does that) even if the probability of landing the punch is in fact higher given the tell compared to some relevant baseline. Since it is not a probability raiser, *a fortiori* this is not a case of a probability raiser that fails to be a means.

We now have two responses to cases that appear to be counterexamples to probability raising as a sufficient condition on being a means: make sure the case concerns actions in the relevant sense and make sure the action helps to make the probability of some end higher. Given these constraints, it is more difficult to generate a counterexample to sufficiency.

Kolodny has a different solution to the boxer's case.<sup>8</sup> He says that a means must be an action such that the probability that the action *helps to bring about the end*, conditional on the action, is positive. And he thinks the boxer's tell does not help to bring about the end of landing a punch. This solution differs from mine at least nominally, for Kolodny has a bringing-about relation between an action

7 Cf. Lin, "Simple Probabilistic Promotion." Lin goes on to characterize this in-virtue-of relation as causal, but I think that is unduly restrictive. Side note: this clarification about probability raisers has interesting applications in decision theory. In Newcomb's paradox, for example, choosing one box arguably does not make the probability that there is \$1 million in the opaque box higher even if the  $\Pr(\$1 \text{ million in opaque box} \mid \text{choose one box})$  is higher than the relevant baseline. Arguably, it is the categorical basis for the tendency to choose one box (detectable by the predictor some time before the present choice) that helps to raise the probability that there is \$1 million in the opaque box. So choosing one box might not be the best means to getting the most money under a probability-raising theory of means.

8 See Kolodny, "Instrumental Reasons."

and an end, whereas I have a making-it-the-case relation between an action and a conditional probability. I think the following case might help bring Kolodny's constraint into focus and show why I think it is mistaken.

Suppose a rogue state has launched a nuclear warhead and the end is to neutralize the threat. A missile defense system is in place, and one intercept has already been fired with .5 probability of neutralizing the threat. Let us assume this .5 probability is the relevant baseline. You have three available actions and you can only do one. You can do nothing, leaving the probability of neutralizing the threat at .5. You can fire another intercept that raises the probability of neutralizing the threat to .75, for if the first misses, there is .5 probability that the second hits. Finally, you can press a button that sends an electromagnetic pulse through the atmosphere that raises the probability of neutralizing the threat to .7, for the signal has .4 probability of deactivating the nuclear warhead before the first intercept comes close, combined with even odds that the first intercept works if the pulse does not.

It seems that you have most derived, means-based reason to fire the second intercept rather than push the button that sends a pulse or do nothing, for it raises the probability of the end from .5 to .75, whereas pressing the button only raises it to .7 and doing nothing leaves it at .5. But notice that, given that you fire the second intercept, it only has a probability of .25 of *helping to bring about* the end, for it can only do that if the first intercept misses, in which case it has even odds of intercepting. By contrast, the probability that pressing the button helps to bring about the end is .4, for it will work (or not) before the first intercept arrives. In light of this, it looks like Kolodny's condition will deliver greater derived reason to press the button than to fire the second intercept even though firing the second intercept makes the end more likely than pressing the button. I think that is the wrong result.

To be fair, Kolodny wants the "helping to bring about" clause to be read broadly, and maybe it can be read broadly enough so that his condition of helping to bring about an end is extensionally equivalent to my condition of *making* the end more probable. If so, I think the probability-raising condition I suggest is a clearer way to identify that extension, and the better candidate for what it is to be a means.

Let me give one more schematic case to support probability raising as sufficient for being a means. Suppose for the sake of argument that we add some extra necessary condition such that to be a means is to be a probability raiser *plus some x-factor*. This would rule out those actions that lack the x-factor even if they are really excellent probability raisers. Now consider a case in which Alice has

two ends,  $e_x$  and  $e_y$ , that are probabilistically independent of one another, and two available actions,  $a_x$  and  $a_y$ . Further:

$a_x$  raises the probability of  $e_x$  by .5, and the x-factor is present.

$a_y$  raises the probability of  $e_x$  by .5, and the x-factor is present.

$a_x$  raises the probability of  $e_y$  by .99, and the x-factor is absent.

$a_y$  raises the probability of  $e_y$  by .5, and the x-factor is present.

If the x-factor is a necessary condition, overall  $a_y$  is a means to *both* ends, while  $a_x$  is a means to only *one* end. If we use probability raising as a guide to the weight of transmitted reasons, Alice would have most reason to do  $a_y$ , for it raises the probability of each end by .5. Meanwhile,  $a_x$  lacks the x-factor vis-à-vis  $e_y$ , so it is not a means to that end and inherits no transmitted reason weight from that end even though it makes that end almost certain to obtain, and far more likely to obtain than does  $a_y$ .

That is the wrong result. If ends are what ultimately matter, it is hard to see why  $a_y$  would derivatively matter more than  $a_x$ . After all,  $a_x$  *makes* the things that ultimately matter more likely to occur compared to  $a_y$ . This is not a conclusive criticism of all x-factor conditions, of course, and there is the possibility of other objections against probability raising as sufficient for being a means. But I hope I have said enough to justify my focus on pure probability-raising accounts. That is the task for the rest of the paper.

## 2. FRAMEWORK

In what follows, I will work within a certain theoretical apparatus. Let us suppose that, for a given agent,  $S$ , and context of choice,  $C$ , there is a finite set of actions,  $A = \{a_0, a_1, \dots, a_n\}$ , where these actions are (i) available to  $S$  in  $C$ , (ii) jointly exhaustive, and (iii) mutually incompatible with one another.<sup>9</sup> I will not spell out what it is to be an available action (though I have said the actions need to involve agential activity over which one exercises some control), or whether there is a privileged set of exhaustive and incompatible actions to consider for a given agent in a given context. Those are debatable issues and plausible theories on those fronts can be plugged into the general framework here.

Suppose also that for an agent and a context there is a unique, finite set of ends,  $E = \{e_0, e_1, \dots, e_n\}$ . And suppose that for any end,  $e_i \in E$ , its ultimate weight can be modeled by assigning it a positive number,  $g_{ei}$ , on some cardinal scale of weights (possibly different scales for different ends).

9 The approach here is consistent with contrastivism about reasons (see Snedegar, “Contrastive Reasons and Promotion” and “Reason Claims and Contrastivism about Reasons”).

Relative to an agent ( $S$ ), context of choice ( $C$ ), set of available actions ( $A$ ) and set of ends ( $E$ ), we want to analyze what it would be for some  $a_i \in A$  to be a means to some  $e_i \in E$  and how much weight transmits from a given end to its means. The point of specifying the analysis relative to some  $S$ ,  $C$ ,  $A$ , and  $E$  is to model as best we can what agents are faced with when they act, and what considerations other parties can draw on for evaluating the actions of others. We want to do so in a way that allows sensible comparisons, trading off and aggregation of what one has reason to do without shifting the relevant context of choice and options to choose from in the process.

I will proceed by considering various *prima facie* plausible baselines for probability raising. An appendix will discuss some related matters, including what it is to transmit a reason *against* an action, and how to approach the nettlesome problems of aggregation to oughts, incommensurability, and defeat.

### 3. REFRAINING AS AN INDEPENDENT ACTION

For an action  $a_i \in A$  at a context of choice  $C$  the obvious baseline to consider is the case of not- $a_i$ . Does  $a_i$  raise the probability of some end relative to not- $a_i$ ? There are a number of ways to regiment the idea. Let us consider them one by one.

First, one might appeal to *refraining from  $a_i$ -ing* as its own independent action in the set of available actions. The view I have in mind would deviate from some standard ways of characterizing what it is to refrain. On one standard view, one can refrain from doing an action *while* doing some other action. If so, I can refrain from staying home at the same time I do a distinct action, like going to the park. Of course, such a view would not allow us to construct a set of actions  $A$  with both *go to the park* and *refrain from staying home* in the set, for even if these are two available actions they are not incompatible, and incompatibility is a constraint on actions in  $A$ . On another standard view, *what it is* to refrain from one action just is to do some incompatible action. So by going to the park I refrain from staying home. Again, this view does not allow us to construct a set of actions with both *go to the park* and *refrain from staying home* in the set, for they are not incompatible.

Set those characterizations of refraining to one side for a moment. I want to consider a nonstandard view in which refraining from  $a$ -ing is an available action that is not only incompatible with  $a$ -ing, but is also incompatible with other actions incompatible with  $a$ -ing. You might doubt that this is the best view of refraining, but let us grant that it is intelligible and let us consider the view that an action is a means to an end just in case it raises the probability of the end relative to refraining so understood. More generally, for a given  $a_i$  the relevant baseline

would be (i) refraining from  $a_i$  if  $a_i$  is *not* a refraining, or (ii) the non-refraining counterpart to  $a_i$  if it is a refraining. If we let “ $r(a_i)$ ” be the symmetric refraining function on  $a_i$ , and use this for the relevant baseline, the proposal is this: for some  $a_i \in A$  to be a means to some  $e_i \in E$  is for:

1.  $a_i \neq e_i$ ,<sup>10</sup> and
2. Partly in virtue of  $a_i$ ,  $\Pr(e_i \mid a_i) > \Pr(e_i \mid r(a_i))$

The means-based weight for  $a_i$ ,  $w_{ai}$ , which is transmitted from the end-based weight for  $e_i$ ,  $g_{ei}$ , would be proportional to the degree to which  $a_i$  raises the probability of  $e_i$  from the baseline as follows:

3.  $w_{ai} = g_{ei} \times (\Pr(e_i \mid a_i) - \Pr(e_i \mid r(a_i)))$

On this analysis, if finishing the paper is among my ends, and the available actions are {stay home, refrain from staying home, go to the park, refrain from going to the park, go to the movies, refrain from going to the movies}, then staying home is a means to finishing this paper just in case that action makes greater the  $\Pr(\text{I finish this paper} \mid \text{I stay home})$  compared to the  $\Pr(\text{I finish this paper} \mid \text{I refrain from staying home})$ . The big problem with this view is that each of staying home and refraining from staying home might make it highly likely that I finish the paper. Suppose, for instance, the following probabilities:

$$\begin{aligned} \Pr(\text{I finish the paper} \mid \text{I stay home}) &= .75 \\ \Pr(\text{I finish the paper} \mid \text{I refrain from staying home}) &= .75 \end{aligned}$$

Either way, I am very likely to finish the paper. And because staying home does not raise the probability of finishing the paper relative to refraining from doing so I would have no means-based reason to stay home if we take refraining to be the relevant baseline. This is particularly odd when we compare some other available actions that make finishing the paper much less likely. For the other actions available to me, suppose the following conditional probabilities:<sup>11</sup>

$$\begin{aligned} \Pr(\text{I finish the paper} \mid \text{I stay home}) &= .75 \\ \Pr(\text{I finish the paper} \mid \text{I refrain from staying home}) &= .75 \\ \Pr(\text{I finish the paper} \mid \text{I go to the park}) &= .5 \\ \Pr(\text{I finish the paper} \mid \text{I refrain from going to the park}) &= .4 \end{aligned}$$

10 An end cannot be a means to itself on pain of double counting the weight of reasons that support it, once as an ultimate reason and once as a transmitted reason.

11 Henceforth, let it be understood that the examples involving higher conditional probabilities compared to a baseline include the requisite relation of *making* the end more probable unless otherwise specified.

If we take refraining as the baseline, my end of finishing the paper gives me no reason to stay home but some reason to go to the park. Staying home would not be a means at all whereas going to the park would. These are the wrong results. When defining means and the weights of reasons supporting them, we do not want the two most successful actions (probability-wise) to be baselines for each other, nor do we want any two poorer actions to be baselines for each other.

#### 4. DOING SOME ALTERNATIVE

To get away from confounding pair-wise comparisons, suppose we take the baseline for any given  $a_i \in A$  to be *doing some alternative to  $a_i$* , in which an alternative to  $a_i$  is an action in the set  $A$  incompatible with  $a_i$  (we saw that this is one view of what it is to refrain from  $a_i$ , and it is also one interpretation of the referent of “not-  $a_i$ ”). On this view, for an action to be a means is for that action to raise the probability of an end relative to *doing something else*, and the amount of reason transmitted is proportional to how much that probability is raised. More formally, for some  $a_i \in A$  to be a means to some  $e_i \in E$  is for:

1.  $a_i \neq e_i$ , and
2. Partly in virtue of  $a_i$ ,  $\Pr(e_i \mid a_i) > (\Pr(e_i \mid a_0) + \Pr(e_i \mid a_1) + \dots + \Pr(e_i \mid a_n) - \Pr(e_i \mid a_i)) \div n$

The means-based weight,  $w_{ai}$ , relative to ultimate reason weight,  $g_{ei}$ , is this:

3.  $w_{ai} = g_{ei} \times (\Pr(e_i \mid a_i) - ((\Pr(e_i \mid a_0) + \Pr(e_i \mid a_1) + \dots + \Pr(e_i \mid a_n) - \Pr(e_i \mid a_i)) \div n))$

This gets away from confounding pair-wise comparisons, so that is good. The problem is that this baseline makes a given  $a$ 's status as a means, and the weight of reason transmitted to it too sensitive to the presence of *better* means. To see why, consider a set of available actions, {stay home, go to the park, go to the movies}, and suppose the following probabilities:

- $\Pr(\text{I finish the paper} \mid \text{I go to the park}) = .5$
- $\Pr(\text{I finish the paper} \mid \text{I go to the movies}) = .25$
- $\Pr(\text{I finish the paper} \mid \text{I stay home}) = .75$
- $\Pr(\text{I finish the paper} \mid \text{I go to the movies or stay home}) = .5$

Under the analysis, going to the park would not be a means to finishing the paper, for it does not raise the probability of finishing the paper relative to doing something else (going to the movies or staying home). That seems strange given that one of the actions available to me makes finishing the paper much less probable.

In addition, when we use doing something else as a baseline, small changes in the best available action *could* turn something that is not a means into a means. To see this, contrast the above example with the following:

- Pr (I finish the paper | I go to the park) = .5  
 Pr (I finish the paper | I go to the movies) = .25  
 Pr (I finish the paper | I stay home) = .74  
 Pr (I finish the paper | I go to the movies or stay home) = .495

In this scenario, going to the park *is* a means with some weight. But if the decrease in Pr (I finish the paper | I go to the movies or stay home) is due to the decrease in Pr (I finish the paper | I stay home), we have a pair of cases in which the status of going to the park as a means depends on how effective my *best option* is at probabilizing the end.<sup>12</sup> Again, that is the wrong result. When we use *doing something else* as the relevant baseline, it accords too much influence to the best available actions.

## 5. BEFORE VS. AFTER

A similar problem afflicts probability-raising accounts in which the relevant baseline is the probability of the end *prior to acting*.<sup>13</sup> Consider a case in which the probability that I finish the paper is very high simply because it is very likely that I will stay home and very likely that I will finish the paper given that I stay home. Indeed, suppose that this makes the probability of finishing the paper much higher than the probability of finishing the paper given that I go to the park. In that case, going to the park is not a means at all, even if it makes the probability of finishing the paper much higher than something else I could do, like going to the movies. That is the wrong result. Using the probability of the end prior to acting as a baseline gives too much influence to my best option and how likely I am to perform it.

## 6. WHAT WOULD OTHERWISE HAPPEN

One might seek a baseline in whatever happens in the nearest possible world where  $a_i$  fails to occur (yet another possible referent of “not- $a_i$ ”). Let us then condition on  $\sim a_i$ , where this operation is understood to take us to the nearest

12 It is not generally the case that  $\Pr(e | (a_1 \text{ or } a_2)) = [\Pr(e | a_1) + \Pr(e | a_2)] \div 2$ . But there are cases where this occurs.

13 Lin (“Simple Probabilistic Promotion”) has this probability-raising theory of promotion. He acknowledges the problem I point out in the text.

possible world(s) where  $a_i$  does not occur, and ask after the probability of  $e_i$  in that/those world(s).<sup>14</sup> On this view, for some  $a_i \in A$  to be a means to some  $e_i \in E$  is for:

1.  $a_i \neq e_i$ , and
2. Partly in virtue of  $a_i$ ,  $\Pr(e_i \mid a_i) > \Pr(e_i \mid \sim a_i)$

The means-based weight,  $w_{ai}$ , relative to ultimate reason weight,  $g_{ei}$ , is:

3.  $w_{ai} = g_{ei} \times (\Pr(e_i \mid a_i) - \Pr(e_i \mid \sim a_i))$

A special case here deserves special attention—the case in which  $\sim a_i$  takes us to a world where one *fails to act at all*; i.e.,  $C$  and  $A$  are held fixed and  $S$  fails to do any action in  $A$ . I think it would be a mistake to appeal to a nonaction for a baseline. Suppose the worst thing I could *do* vis-à-vis finishing the paper is to go to the movies, where  $\Pr(\text{I finish this paper} \mid \text{I go to the movies}) = .25$ , as before, and *no other available action makes finishing the paper less likely*. It just so happens that I will go to the movies, but if I were not to go to the movies, I would pass out, where  $\Pr(\text{I finish this paper} \mid \text{I pass out}) = .15$  (I might be able to finish it after I come to, but I will be distracted by medical concerns). It follows on the present analysis that going to the movies is a means to finishing the paper. For it probabilizes the end relative to what happens in the nearest world in which I do not go to the movies.

That might seem counterintuitive by itself. How could the worst thing I could do be a means to my end? A deeper problem is that allowing for nonaction baselines frustrates the first-personal goal of helping to decide what to do and the second- and third-personal goals of evaluating what *is done* against what *could have been done*. I cannot choose an action over a nonaction—I cannot decide to go to the movies rather than pass out. And it seems illegitimate to take into account nonactions when evaluating what is done. It is illegitimate, for instance, to positively evaluate going to the movies to some degree because of the possibility that I pass out. So I think it is inappropriate to go outside the set of available actions to set baselines.

The point might be easier to see with a case in which nonaction makes achieving an end most likely. So consider a context in which my end is that my mother is taken care of. I could do a lot in this regard. I could try to take care of her myself, I could make various arrangements, etc. Other things would frustrate this end, like going to the casino. But suppose that, for each available action, it turns

14 Finlay (“The Reasons that Matter,” 8n19) arguably has this view. In *Confusion of Tongues*, 38–46, he seems to have a different view.

out that in the nearest possible world in which it does not occur I pass out.<sup>15</sup> It also turns out that in those nearest possible worlds others will certainly ensure my mother is taken care of. In this case, every action I could do actually reduces the probability of the end relative to the  $\sim a_i$  baseline, so none of them are means. That is, for all  $a_i \in A$ ,  $\Pr(e_i \mid a_i) < \Pr(e_i \mid \sim a_i)$ . On such a view, there is no reason to choose an action in virtue of being a means to my end. But surely some of the things I can do better achieve my end than others—arranging for my mother’s care is better than going to the casino. To extend the point made above, I cannot use the results of the analysis to make any decisions, and those wishing to evaluate my conduct cannot appeal to the fact that things go best when I pass out.

It is tempting to say that the best thing to do is that which will bring about my passing out. But that is a very particular kind of case. The hard case to focus on is one where nothing I can do probabilizes passing out, and each thing I can do fails to probabilize the end relative to what happens if I do not do that thing (where we stipulate that, for each action, if I do not do it I pass out). This kind of case helps show why *not acting* cannot be a relevant baseline.

All this can be avoided by allowing “ $\sim a_i$ ” to range only over the nearest possible world(s) in which one *performs an action* other than  $a_i$ .<sup>16</sup> This is only a partial fix. It might help meet the deliberative and evaluative concerns, but it still misidentifies means. If in the nearest possible world in which I perform an action other than  $a_i$  I perform the action that makes the end most probable, this should not thereby make  $a_i$  a non-means. It should be highly relevant whether I can do things other than  $a_i$  that make my end even less likely than does  $a_i$ . As with using the performance of some alternative action as a baseline, this one is too sensitive to better means and not sensitive enough to worse ones.

Note also that, like the refraining function above, subjunctives do not pick out a unique baseline for all actions in  $A$ . There is therefore potential to assign means-based weights in a way that skews the comparison of reasons for the actions in  $A$ . In the paper-writing case, suppose a set of available actions  $A = \{\text{stay home, go to the coffee shop, go to the park, go to the movies}\}$ , in which going to the coffee shop and staying home make finishing the paper equally likely and more likely than anything else I could do. More specifically, suppose:

15 This seems to entail that I will actually pass out, in which case you might think that the actions in  $A$  are not available after all. But we should not take what actually happens to limit what can happen.

16 This is similar to the baseline Jackson and Pargetter fix when the question is whether to do  $a_i$  (“Oughts, Options, and Actualism,” 246), although, in effect, they think this context of choice identifies the set of available actions with the set  $\{a_i, \sim a_i\}$ . I will revisit this issue below.

- $\Pr(I \text{ finish the paper} \mid I \text{ stay home}) = .75$   
 $\Pr(I \text{ finish the paper} \mid I \text{ go to the coffee shop}) = .75$   
 $\Pr(I \text{ finish the paper} \mid I \text{ go to the park}) = .5$   
 $\Pr(I \text{ finish the paper} \mid I \text{ go to the movies}) = .25$

Further, I *will* go to the coffee shop. If I do not go to the coffee shop, I go to the movies. If I do not stay home, I go to the coffee shop. On this fact pattern, staying home is not even a means, and no reason transmits to it, while going to the coffee shop is a very good means, supported by weighty means-based reason. That is the wrong result.<sup>17</sup> The problem is a result of shifting baselines.

## 7. THE STATUS QUO

Could the baseline be the *status quo*, such that an action is a means to an end if it raises the probability of that end relative to the status quo?<sup>18</sup> Though it is not entirely clear what this baseline amounts to, it sounds like the analysis would make it impossible to have among *A* the action of *carrying on with the status quo*, and for that action to be a good means to one's end. For carrying on with the status quo would not raise the probability of an end relative to the status quo. That is the wrong result. Sometimes we are already in the process of taking a means to one of our ends. Indeed, sometimes we are already taking the best means and we should carry on.<sup>19</sup>

## 8. ANY POSITIVE CONDITIONAL PROBABILITY

Now consider the zero baseline: to be a means to an end just is to make that end probable *to any positive degree*. More formally, for some  $a_i \in A$  to be a means to some  $e_i \in E$  is for:

- 17 Cf. Behrends and DiPaolo, "Finlay and Schroeder on Promoting a Desire," 2–3  
 18 Mark Schroeder defends the view that one has reason to do that which *promotes* the object of any of one's desires, where promotion is understood probabilistically but where the weight of reason is not proportional to the degree of promotion (*Slaves of the Passions*, 113). He characterizes the baseline for promotion both in terms of *doing nothing* and the *status quo*. To be fair, his primary aim in that work is not to settle the details of the promotion relation, but to press a defensive strategy for Humean theories of reasons.  
 19 Cf. Evers, "Humean Agent-Neutral Reasons?" 60; and Behrends and DiPaolo, "Finlay and Schroeder on Promoting a Desire," 4–5. In their criticism, Evers and Behrends and DiPaolo focus on the "do nothing" baseline, while in this section I focus on the "status quo" baseline. It is not at all clear that the two baselines are equivalent, for the status quo might not involve doing nothing.

1.  $a_i \neq e_i$ , and
2. Partly in virtue of  $a_i$ ,  $\Pr(e_i \mid a_i) > 0$

The means-based weight,  $w_{ai}$ , relative to ultimate reason weight,  $g_{ei}$ , is:

3.  $w_{ai} = g_{ei} \times \Pr(e_i \mid a_i)$

This analysis is nice and simple.<sup>20</sup> By including the “in virtue of” clause in 2, we avoid some very natural concerns. It is natural to wonder, for example, about the special case in which  $\Pr(e_i) = 1$ . Does this analysis make all actions ( $\neq e_i$ ) in any set  $A$  means, supported by as much reason as the end itself? No, thanks to the “in virtue of” clause. If the  $\Pr(e_i) = 1$  *because* of some action or actions in  $A$ , then those actions that are making the probability go to 1 will be probability raisers and they will get some weight transmitted to them. But if no action of yours is helping to make it the case that  $\Pr(e_i) = 1$ , then they will not count as means and will receive no weight from the end. Whatever the probability of  $e_i$  is, the question for this analysis is whether some action of yours would make its probability some amount greater than zero. Then and only then is it a probability raiser.

The troublesome case for this proposal is a case in which some action of yours makes the probability of some end greater than zero, but the action makes the end *as unlikely as possible*. If  $a_i$  is the absolute worst thing I could do vis-à-vis  $e_i$  (probabilistically speaking), but doing so still yields a positive probability for  $e_i$ , the analysis counts  $a_i$  as a means to  $e_i$ , supported by some weight. Most pointedly,  $a_i$  would be a means supported by some weight even if every other action available to you makes the end much more likely. Suppose, for example, that the prior probability of finishing the paper is .5, and of my three available actions (stay home, go to the park, go to the movies), going to the movies drops the probability to .1, while going to the park keeps it at .5 and staying home raises it to .9. This is a case in which going to the movies is a difference maker, probability wise, but in a bad way. Calling it a means seems inapt—if I make finishing the paper as unlikely as I can, I am not taking a means to finishing it. To solve this concern, some relevant baseline other than zero is called for.<sup>21</sup>

20 And I think it is the one used by Kolodny (“Instrumental Reasons”). This is as good a place as any to point out that Kolodny needs a clause in his transmission principle to block iterations of the principle that would generate reasons that do not flow from ultimate ends. We handle the problem by defining means and analyzing weights directly in terms of those things we have *ultimate* reason to do.

21 As one reviewer pointed out, one advantage of the zero baseline is that it makes the status as a means independent of the set of available actions,  $A$ , that we construct. By contrast, my favored view below is highly “menu dependent”: whether an action is a means and how much reason gets transmitted to it depends on the set of actions,  $A$ , in which it is found, though I

## 9. THE WORST ONE COULD DO

To sum up the lessons from above, we are in need of a unique baseline for all actions in  $A$ , one that is itself an action and that takes the availability of worse actions vis-à-vis an end to be more probative than the availability of better actions, but which does not count any action with positive conditional probability for an end as a means. So let us consider the view that an action is a means to your end insofar as it raises the probability of the end relative to the worst you could do. More formally, let us characterize the worst you can do like this: for a given  $e_i \in E$ , the worst you can do is that  $a_u \in A$  such that  $\Pr(e_i | a_u) \leq \Pr(e_i | a_j)$ , for all  $a_j \in A$ . Then for some  $a_i \in A$  to be a means to some  $e_i \in E$  is for:

1.  $a_i \neq e_i$ , and
2. Partly in virtue of  $a_i$ ,  $\Pr(e_i | a_i) > \Pr(e_i | a_u)$

The means-based weight,  $w_{ai}$ , relative to ultimate reason weight,  $g_{ei}$ , is:

$$3. w_{ai} = g_{ei} \times (\Pr(e_i | a_i) - \Pr(e_i | a_u))$$

I think this is the best probability-raising account of means and the weight of reasons supporting them. It meets the desiderata spelled out above. It is an action, it makes the weight of reasons to take means insensitive to better means and it does not necessarily count actions with any positive conditional probability as means.

The present proposal has some similarities with actualism, defended by Jackson and Pargetter ("Oughts, Options, and Actualism"). To compare the views, let us consider the paper-writing case again, where  $A$  is {stay home, go to the park, go to the movies}, and suppose that if I go to the park I will get injured, spend some time in the hospital and likely finish my paper there, whereas if I stay home I stand a good chance of playing video games instead of finishing the paper. The probabilities are:

$$\begin{aligned} \Pr(\text{I finish the paper} \mid \text{I stay home}) &= .5 \\ \Pr(\text{I finish the paper} \mid \text{I go to the park}) &= .75 \\ \Pr(\text{I finish the paper} \mid \text{I go to the movies}) &= .25 \end{aligned}$$

On the preferred probability-raising view, given  $C$  and  $A$ , I have more derivative reason to go to the park than to stay home. This is so even if  $\Pr(\text{I finish the paper} \mid \text{I stay home and work on the paper}) = 1$ .

This is in one respect similar to actualism, where the question concerns what

---

leave open whether there are privileged sets. In any event, for these reasons I am conflicted about my favored analysis. Positive conditional probability is a close second to the preferred analysis.

one ought to do, and the actualist answers: do that action where things turn out best. Two qualifications characteristic of actualism are (a) the value of an option (available action) is fixed by what actually happens should you choose that option, where one treats whatever happens after doing the action as an exogenous variable, even those future actions that will be available to the relevant agent, and (b) when the question is whether to do some particular action, the answer turns on whether the value of the outcome were one to do that action is greater than the value of the outcome were one to fail to do that action (go to the nearest possible world(s) to check).

The probability-raising analysis has an analogue to (a). Everything outside the set of alternatives at a context, including future actions that are not available to *S* at *C*, is treated as an exogenous variable. Though the  $\text{Pr}(\text{I finish the paper} \mid \text{I stay home and work on the paper}) = 1$ , and though I might have most derived reason to stay home and work on the paper when considering an *A* that includes this action string, this does not mean I have most reason to stay home. Some object to this, but the problem can be handled in the usual way. Though it is true that under the set of alternatives considered above one has more means-based reason to go to the park than to stay home, in a different context with a set of actions that includes *staying home and working on the paper*, one might well have more means-based reason to do that than to do any alternative that involves going to the park. Indeed, it is open to theorists to argue for some privileged context of choice, or a privileged set *A* in a given context, that is most appropriate for deliberation or evaluation. And the privileged *A* might include strings of actions like *staying at home and working on the paper*. Even if there is no privileged *A*, there would still be constructible *As* that include these strings of actions and folks who think one has most reason to stay home might have in mind not just staying home, but the string *stay home and work on the paper*.

As for (b), the probability-raising analysis on offer has no analogue. The only sets of alternatives are exhaustive and mutually incompatible ones, and we earlier rejected any baseline fixed by subjunctives that deliver pairwise comparisons, where the comparison class can shift depending on what action we are assessing. The probability-raising analysis thereby avoids some of the counterintuitive consequences of actualism, such as holding that one ought to do a terrible action when one would otherwise do something *even worse*, and despite the fact that much better options are also available. The probability-raising analysis always keeps in focus some set of available actions that is exhaustive, and so it will not lose sight of the best options available for bringing about a given end.<sup>22</sup>

22 Of course, the conditional probabilities for actions in the exhaustive set of alternatives will be informed by the probable future actions of the agent, which might be heinous. Worries

## 10. CONCLUSION

Given some conception of ends as things supported by ultimate reasons, I have tried to shed light on what it is to be a means to those ends, and how much reason is transmitted from an end to its means. The study yields a probability-raising theory, where an action is a means to an end relative to the baseline of the worst action available (probability wise). I think this is the best pure probability-raising theory, and so it is the one to compare with any alternative to the pure probability-raising approach.

## APPENDIX: ANTI-ENDS AND AGGREGATION

The present theory can be extended to deal with negative reason locutions, such as “you have some reason not to  $\phi$ ,” and with the problem of aggregation. Let me start with the first. It surely is intelligible to speak of reasons not to do things, as when one says, “You have a reason not to [/to not] kill innocents.” Alternatively, we might speak of a reason *against* acting, as in, “You have a reason against killing innocents.” These locutions raise some interesting issues. For example, when we speak of a reason for not-killing we need to know what “not” does *qua* action modifier. And when we speak of a reason *against killing* we need to know what reasons *against* action are and how they differ from reasons *for* action.

Let me focus on the first locution, and what “not” might be doing as an action modifier. We have already seen candidates for the referent of a “negated-action” clause—those considered in the discussion of baselines. For some  $a_i$ , “not- $a_i$ ” could refer to:

- (1)  $r(a_i)$ —we assume here the nonstandard view of refraining in section 3,
- (2) doing something in  $A$  other than  $a_i$ , or
- (3) whatever occurs in the nearest possible world where  $a_i$  does not.

There are problems with each proposal. According to (1), “You have a reason not to [/to not] kill innocents” would refer to a reason for refraining from killing innocents. You act in accordance with this reason so long as you refrain from killing innocents. The problem here is that this refraining is typically one action

---

about this will be met as above, where we can restrict our attention to certain contexts of choice and sets of actions. Indicative conditionals complicate matters. If I say “If/given that I will not work on the paper here at home, I have most reason to go to the coffee shop” the antecedent is influencing the conditional probability assessments by restricting the possible worlds over which the probabilities are determined. If we calculate  $\Pr(\text{I finish the paper} \mid \text{I stay home})$  over worlds in which I do not write the paper at home, it might well be less than  $\Pr(\text{I finish the paper} \mid \text{I stay home})$  calculated over a wider set of possible worlds.

among many available ones. So suppose the relevant  $A = \{\text{kill innocents, refrain from doing so, buy some ice cream}\}$ . Again, we assume that refraining from killing innocents is incompatible with buying ice cream. This might seem wrong, for by having ice cream you do not kill, and you might think this is enough to thereby refrain from killing. That is right under some standard conceptions of what it is to refrain, which can be handled under option (2). However, to see all the options clearly we are working within a nonstandard conception of refraining in (1). For us, doing some action that precludes killing does not entail that one refrains from killing. Refraining from killing can be incompatible with other actions that entail that one does not kill.

Now, you have very weighty reason not to kill innocents, and some reason to buy some ice cream. Intuitively, you should buy ice cream, for that is what you have most reason to do. But if “weighty reason not to kill innocents” refers to weighty reasons *for* refraining from killing them, you would have most reason to refrain from killing innocents, for presumably the reason not to kill innocents is far weightier than your reason to buy ice cream. No ice cream for you! That is sad and wrong. For by buying ice cream you also avoid killing innocents—you can get that result without refraining from killing innocents, considered as an independent action incompatible with the other options.

It looks like (2) fixes this problem, in part by dropping the nonstandard view of refraining. According to it, “reason not to kill innocents” in our ice cream example would refer to reason *for* (refraining from killing innocents or buying ice cream). It is not entirely clear what this amounts to, for one can do either of these alternatives to killing innocents, and we want to know how much reason we have for *each* one, not for doing one or the other. One way forward is to let the strength of the reason not to kill innocents to be a reason of that strength to refrain from killing and a reason of that strength for buying ice cream. This would generate the right results for the case. You would have great reason for refraining from killing innocents and great reason for buying ice cream, plus some additional reason to buy ice cream, presumably based on its tastiness. On this view, you have most reason to buy ice cream, which is a happy result.

However, (2) cannot account for the asymmetry between tragic dilemmas and delightful dilemmas. Think of a version of Williams’s Jim and the Indians case, where Jim must (a) accept the invitation and kill one person or (b) decline the invitation to kill one person, which will lead to the death of twenty by the hand of a chief. Intuitively, there are only negative things to say about the options; that is, there are reasons *against* each one. That is what makes the choice so difficult. But under (2) we are really thinking of reasons for each action—what is naturally thought of as a reason not to do one of the options, or against one

of the options, is really a reason for doing some alternative. What is naturally thought of as a reason not to decline the invitation, for instance, is really a reason to accept and kill.

By itself, this is a little odd. But it is particularly odd when we compare the case with a delightful dilemma. In a delightful dilemma, we might imagine these available actions: (a) out of a group of deserving people, pick one person to be given a genie's wish, or (b) decline to pick, in which case the genie will choose someone. If the only reasons in play are reasons to benefit others, intuitively there are only positive things to say about the options—that is, there are reasons *for* each one. That is what makes this choice difficult. What makes it delightful, and so structurally different from its tragic cousin, is that there are good reasons for each available action, rather than against each option. If (2) is correct, however, there is no such structural difference, for both cases feature reasons *for* each available action. Thus, (2) does not respect the structural difference between tragic and delightful dilemmas.

This can be put in terms of a difference in the *valence* of reasons, where a negative valence toward one action is not the same as a positive valence toward its alternatives. In other words, we should avoid interpreting these valences like vectors and their magnitudes. A vector force of negative magnitude,  $-m$ , in one direction is equivalent to a vector force of  $m$  in the opposite direction. By contrast, I am suggesting that a reason against an action of strength  $-w$  is not equivalent to a reason of strength  $w$  for “opposing” actions.<sup>23</sup>

Last, (3) fails for similar reasons as (1) and (2). It has the problems of (1) insofar as it identifies a single probabilistically open alternative to receive all the weight of a reason not to  $a_i$ . There might be other probabilistically open alternatives other than the one of the nearest possible world where  $a_i$  is absent, and reasons to do these should not be outweighed by a weighty reason that favors the nearest possible world where  $a_i$  is absent. It has the problems of (2) insofar as it fails to respect a structural difference between delightful and tragic dilemmas. Through the lens of (3), all reasons can be interpreted as positive reasons for doing things or for states of affairs.

In the end, I think we should interpret the locution “reason for not  $a_i$ -ing” as referring to a reason *against*  $a_i$ -ing, where a reason against is a reason with negative weight not to be assimilated with a reason of positive weight for  $a_i$ 's

23 Cf. Greenspan (“Asymmetrical Practical Reasons”), who uses a structural difference like the one defended here to think about how different reasons relate to requirements. I understand that Snedegar (“Reasons for and Reasons Against”) also has a paper that discusses these issues, but I have discovered this paper only recently and have not had time to fully consider and benefit from it.

alternatives. We can get negative weights out of a probability-raising analysis of means by positing anti-ends—those things one has ultimate *negative* reason to do (/ultimate reason *against*). Anti-ends transmit their negative weight to their means in the same fashion as ends transmit positive weight to their means. In the big picture,  $E$  includes both ends and anti-ends, where for some  $A$ , each end in  $E$  transmits some positive weight to its means in  $A$ , and each anti-end in  $E$  transmits some negative weight to its means in  $A$ .

With the resources of this paper it would be nice to provide an account of aggregation into oughts. Though that must be left for another day, here are some considerations to bear in mind. First, there is the vexing case where a set of ends  $E$  has a plurality of members, and where the conditional probabilities of some ends are not probabilistically independent of one another, and particularly not independent given some means. I am thinking of cases where the  $\Pr(e_1 \text{ and } e_2 \mid a_i)$  is not calculable from  $\Pr(e_1 \mid a_i)$  and  $\Pr(e_2 \mid a_i)$ . If the action is staying home, for example, it matters whether  $e_1$  is finishing the paper and  $e_2$  is grading student papers (likely not compossible), or whether  $e_1$  is having dinner and  $e_2$  is listening to music (likely compossible and perhaps co-probable). Even if each end is equally important and even if staying home raises the probability of each of these ends individually by .5, staying home might raise the  $\Pr$  (finish paper *and* listen to music) more than the  $\Pr$  (finish the paper *and* grade papers). For this reason, we might not be able to aggregate our derived reasons into oughts simply by adding up the weights transmitted from each end individually to each of its means individually.

Second, we might want aggregation to make room for incommensurability and undercutting defeat. To do so, we might not assume that each  $e_i \in E$  can be assigned weight on the same scale. Some ends, you might think, have a more significant kind of weight than others, and some anti-ends transmit reasons against that should be thought of more like side constraints than weights to be traded off with other weights. One then needs to keep track of which kinds of weights get transmitted to means, where any  $a_i \in A$  might be supported to different degrees by different kinds of weight. Again, this prevents a simple additive approach. But there are other options. One could lexically order the scales of weight so that one ought to do that  $a_i \in A$  that has maximal weight on the highest-ranked scale in which any action registers. Or one could identify scales of negative weight that are most grievous, exclude any action with significant weight in the grievous scales and then maximize the positive weight of the remaining actions. Last, one could take a weighted average of the weights registered in all scales, giving the greatest weights to the highest-ranked scales of positive value and the most grievous scales of negative value. On this view, one can first aggregate within

each scale once the means-based reasons are identified, and create a more complicated aggregation function for weights on different scales, effectively placing them on the same scale after all.

More complicated aggregation functions are surely available to handle the most refined normative palettes. Importantly, given the modeling here, any incommensurability, defeat, enabling, or the like is to be handled either at the level of fixing ends at a context, or in the function from ultimate and transmitted reasons at a context to oughts. All this leaves open the possibility that shifts in *C* can engender shifts not only in the set of ends/anti-ends and the weights of ultimate reasons for/against them, but also shifts in how the aggregation function moves from reasons to oughts.<sup>24</sup>

University of British Columbia  
mbedke@mail.ubc.ca

#### REFERENCES

- Bedke, Matthew S. "The Iffiest Oughts." *Ethics* 119, no. 4 (July 2009): 672–98.
- Behrends, Jeff, and Joshua DiPaolo. "Finlay and Schroeder on Promoting a Desire." *Journal of Ethics and Social Philosophy* (December 2011).
- . "Probabilistic Promotion Revisited." *Philosophical Studies* 173, no. 7 (July 2015): 1735–54.
- Bratman, Michael E. "Intention, Practical Rationality, and Self-Governance." *Ethics* 119, no. 3 (April 2009): 411–43.
- Coates, D. Justin. "An Actual-Sequence Theory of Promotion." *Journal of Ethics and Social Philosophy* (January 2014).
- Darwall, Stephen L. *Impartial Reason*. Ithaca: Cornell University Press, 1983.
- Evers, Daan. "Humean Agent-Neutral Reasons?" *Philosophical Explorations* 12, no. 1 (2009): 55–67.
- Finlay, Stephen. *Confusion of Tongues*. Oxford: Oxford University Press, 2014.
- . "The Reasons that Matter." *Australasian Journal of Philosophy* 84, no. 4 (2006): 1–20.
- Greenspan, Patricia. "Asymmetrical Practical Reasons." In *Experience and Analysis*, edited by Maria E. Reicher and Johan C. Marek, 387–94. Vienna: ÖBV & HPT, 2005.

<sup>24</sup> I would like to thank Niko Kolodny, Ryan Millsap, and Jacob Stegenga for great conversations that helped this paper along. I also thank several anonymous referees for helpful comments on previous drafts.

- Hitchcock, Christopher. "Do All and Only Causes Raise the Probabilities of Effects?" In *Causation and Counterfactuals*, edited by John Collins, Ned Hall, and L. A. Paul, 403–18. Cambridge, MA: MIT Press 2004.
- Horty, John F. "Reasons as Defaults." *Philosophers' Imprint* 7, no. 3 (April 2007): 1–28.
- Jackson, Frank, and Robert Pargetter. "Oughts, Options, and Actualism." *The Philosophical Review* 95, no. 2 (April 1986): 233–55.
- Kolodny, Niko. "Instrumental Reasons." In *The Oxford Handbook of Reasons and Normativity*, edited by Daniel Star. Oxford: Oxford University Press, forthcoming.
- Lin, Eden. "Simple Probabilistic Promotion." *Philosophical and Phenomenological Research*, forthcoming.
- Raz, Joseph. "Instrumental Rationality: A Reprise." *Journal of Ethics and Social Philosophy* (December 2005).
- . "The Myth of Instrumental Rationality." *Journal of Ethics and Social Philosophy* 1, no. 1 (April 2005).
- Schroeder, Mark. "Means-End Coherence, Stringency, and Subjective Reasons." *Philosophical Studies* 143, no. 2 (March 2009): 223–48.
- . *Slaves of the Passions*. Oxford: Oxford University Press, 2007.
- Sharadin, Nathaniel. "Checking the Neighborhood: A Reply to DiPaolo and Behrends on Promotion." *Journal of Ethics and Social Philosophy* (February 2016).
- . (2015) "Problems for Pure Probabilism about Promotion (and a Disjunctive Alternative)." *Philosophical Studies* 172, no. 5 (May 2015): 1371–86.
- Snedegar, Justin. "Contrastive Reasons and Promotion." *Ethics* 125, no. 1 (October 2014): 39–63.
- . "Reason Claims and Contrastivism about Reasons." *Philosophical Studies* 166, no. 2 (November 2013): 231–42.
- . "Reasons for and Reasons Against." *Philosophical Studies* (forthcoming).
- Stegenga, Jacob. "Probabilizing the End." *Philosophical Studies* 165, no. 1 (August 2013): 95–112.
- Suppes, Patrick. *A Probabilistic Theory of Causality*. Amsterdam: North-Holland Publishing Company, 1970.