

# Deontological Decision Theory and Agent-Centered Options\*

*Seth Lazar*

---

Deontologists have long been upbraided for lacking an account of justified decision-making under risk and uncertainty. One response is to develop a deontological decision theory—a set of necessary and sufficient conditions for an act’s being permissible given an agent’s imperfect information. In this article, I show that deontologists can make more use of regular decision theory than some might have **thought** but that we must adapt decision theory to accommodate agent-centered options—permissions to favor or sacrifice our own interests, when doing so is overall morally worse. Accommodating options requires more than just amending the decision-theoretic ‘value function’. We must change the decision rule as well.

---

## I. INTRODUCTION

If moral theory is silent except when all the facts are known, then it is, for all practical purposes, mute. All action is shrouded in uncertainty. We can deal with our doubt in different ways. One of the most promising is to articulate and defend a *criterion of subjective permissibility*—a set of

\* This article was presented at the workshop on Ethics, Uncertainty and Decision Theory at Australian National University, July 2015, and at Arizona State, University of Southern California, and University of North Carolina, Greensboro. Thanks to my hosts and to the participants for their comments. Particular thanks for reading drafts to Christian Barry, Garrett Cullity, David Enoch, Peter A. Graham, Alan Hájek, Caspar Hare, Josef Holden, Frank Jackson, Alex King, Avery Kolers, Chad Lee-Stronach, Terrance McConnell, Doug Portmore, Jon Quong, Kai Spiekermann, Sharon Street, Larry Temkin, Laura Valentini, Bas Van der Vossen, Robbie Williams, and Michael Zimmerman. The idea for the article came out of discussions with Andrew Sepielli, Tom Hurka, and in particular Sergio Tenenbaum, after a visit to Toronto. Thanks to them all, as well as to the Editor and the referees for their helpful **comments**.

*Ethics* 127 (April 2017): 1–31

© 2017 by The University of Chicago. All rights reserved. 0014-1704/2017/12703-00XX\$10.00

necessary and sufficient conditions for an act's being permissible in light of an agent's uncertainty about the nonmoral facts.<sup>1</sup>

Deontology is a broad church. But on one influential view, which seems to me nearly right, deontologists deny that what we ought to do is determined by how good or bad the states of affairs are that we realize through our actions.<sup>2</sup> They believe that sometimes it is permissible, and sometimes required, to bring about suboptimal states of affairs. In slogan form: the right is prior to the good.

Deontologists struggle with uncertainty.<sup>3</sup> Although they have addressed the topic piecemeal, none has ever defended a criterion of subjective permissibility. Some consequentialists have offered them **one** but typically only to set them up for reductio.<sup>4</sup> My project is to answer this

1. Some terminological points: subjective permissibility is an umbrella term including different epistemic standards, such as belief-relative, reasonable-belief-relative, and evidence-relative permissibility. I take no stand in this article on the important question of which is most appropriate. I also focus exclusively on nonmoral uncertainty. Moral uncertainty is an interesting and related topic, but it is not my focus here. Finally, where decision theorists reserve 'uncertainty' for situations in which we cannot assign probabilities, I use it in the **ordinary language** sense, to cover all cases in which we act under imperfect information. In fact, I consider only cases in which we can assign probabilities (what decision theorists call 'decision-making under risk'). I discuss some implications of our probabilities being imprecise in Sec. VII.

2. Samuel Scheffler, *The Rejection of Consequentialism* (Oxford: Clarendon, 1994).

3. A failing for which they are often upbraided: Shelly Kagan, *The Limits of Morality* (Oxford: Clarendon, 1989), 90ff.; Barbara H. Fried, "What Does Matter? The Case for Killing the Trolley Problem (or Letting It Die)," *Philosophical Quarterly* 62 (2012): 505–29; Elizabeth Ashford, "The Demandingness of Scanlon's Contractualism," *Ethics* 113 (2003): 273–302; Frank Jackson and Michael Smith, "Absolutist Moral Theories and Uncertainty," *Journal of Philosophy* 103 (2006): 267–83; David Sobel, "Backing Away from Libertarian Self-Ownership," *Ethics* 123 (2012): 32–60; Dennis McKerlie, "Rights and Risk," *Canadian Journal of Philosophy* 16 (1986): 239–51.

4. Graham Oddie and Peter Milne, "Act and Value: Expectation and the Representability of Moral Theories," *Theoria* 57 (1991): 42–76; Jackson and Smith, "Absolutist"; Frank Jackson and Michael Smith, "The Implementation Problem for Deontology," in *Weighing Reasons*, ed. Barry Maguire and Errol Lord (Oxford: Oxford University Press, 2015), 279–91; Michael Huemer, "Lexical Priority and the Problem of Risk," *Pacific Philosophical Quarterly* 91 (2010): 332–51; Jennie Louise, "Relativity of Value and the Consequentialist Umbrella," *Philosophical Quarterly* 54 (2004): 518–36; Yoav Isaacs, "Duty and Knowledge," *Philosophical Perspectives* 28 (2014): 95–110. Colyvan et al. are more sympathetic, but they set aside the central deontological tenet that it is permissible to act suboptimally (Mark Colyvan, Damian Cox, and Katie Steele, "Modelling the Moral Dimension of Decisions," *Noûs* 44 [2010]: 503–29, n. 8). Broome, although sympathetic, is less concerned with deontological decision theory than with whether objective moral theories can be represented as teleological (John Broome, *Weighing Goods: Equality, Uncertainty and Time* [Oxford: Blackwell, 1991]). Smith is also more sympathetic, although she does raise a puzzle for subjectivist deontological ethics that she leaves unsolved (Holly M. Smith, "The Subjective Moral Duty to Inform Oneself before Acting," *Ethics* 125 (2014): 11–38. Deontologists do need to answer Smith's **challenge** but should do so, I think, when defending a specific epistemic standard for deontological decision theory, a task that I defer for another occasion.

challenge: to develop a deontological decision theory that delivers plausible verdicts on what we ought to do given our uncertainty and is both compatible with, and grounded in, basic features of deontological ethics. Perhaps surprisingly, I think the right place to start is a standard more naturally associated with the opposing camp: acts are subjectively permissible just in case they maximize expected utility.

Let me start with some rebranding. Although the concept of ‘utility’ in decision theory is malleable enough to mean whatever we want, it is likely to put deontologists off. So let us sacrifice euphony for the sake of clarity and instead refer to ‘choiceworthiness’, which I will define below.<sup>5</sup> In Section II, I suggest that the principle Maximize Expected Choiceworthiness (MEC) is more readily available to deontologists than they and their critics believe. I do not try to argue for MEC but rather to show that it is a genuine contender for the role of deontological decision theory.<sup>6</sup> Why? Because decision theorists have explored decision-making under uncertainty in minute detail. It would be a great benefit to deontologists if they could take advantage of this expertise. What’s more, it would be welcome if moral and rational decision-making under uncertainty could be addressed in similar ways.

And yet, deontologists’ suspicions of MEC are not wholly baseless. In Section III, I show that it is not the standard we are looking for. It describes a sufficient but not necessary condition for subjective permissibility, because it is sometimes subjectively permissible to bring about a sub-optimal outcome. In brief: sometimes an agent is certain that an option that will be very costly to her is objectively permissible, while believing that the less costly alternative risks objective wrongdoing. In such cases, MEC will entail that the more costly option is subjectively required. Since this is clearly counterintuitive in many cases (I offer one that seems decisive), MEC has to go.

Plausible objective deontological theories cannot be represented with a single choiceworthiness function that, combined with the injunction to maximize, can preserve tenable conclusions about subjective permissibility.<sup>7</sup> Some might think that the problem is with **maximizing** and propose an alternative: Satisfice Expected Choiceworthiness.<sup>8</sup> Provided

5. ‘Deontic value’ is another alternative (Smith, “Subjective Moral Duty to Inform”). I am persuaded, however, that understanding agent-relative reasons in terms of value is a mistake (Mark Schroeder, “Teleology, Agent-Relative Value, and ‘Good,’” *Ethics* 117 (2007): 265–95.

6. Thus far I am a fellow traveler with Colyvan, Cox, and Steele (“Modelling”), although we proceed differently—they identify moral constraints on admissible individual utility functions; I speak of a moral choiceworthiness function.

7. Contra Colyvan, Cox, and Steele, “Modelling”; Oddie and Milne, “Act and Value”; among others.

8. See, e.g., Michael Byron, ed., *Satisficing and Maximizing: Moral Theorists on Practical Reason* (Cambridge: Cambridge University Press, 2004).

my act realizes enough of what matters, then it is permissible. Decision theorists disturbed by the cases in Section III are likely to reach for this solution. In Section IV I argue against it.

The shift to satisficing attempts a technical fix for a substantive moral problem. This involves chasing the deontic facts, reshaping the principle to accommodate your intuitions, losing its capacity to guide your judgments in hard cases. Better, instead, to go back to basics, and ask why deontologists affirm options to act suboptimally, before reconstructing our criterion of subjective permissibility. Justifying options is, of course, a task for another paper.<sup>9</sup> But in Section V, I invoke a compelling deontological idea to explain them. It starts with moral status. Having moral status means that your interests have value—they can determine what you and others ought to do. But it also means that you are an end in yourself, not a mere site for the realization of value. This in turn means that you are not required to sacrifice your interests just in case doing so realizes more value overall, nor are you required to serve your own interests, if the alternatives are suboptimal. I will call these *agent-centered options*.

In Section VI I introduce and explain a principle that retains the most useful features of MEC, while capturing the idea that we are not mere sites for the realization of value but instead have self-favoring and self-sacrificing options:

COST: An option is subjectively permissible if and only if, and because, either (a) no other option with greater expected choiceworthiness has reasonable marginal expected costs to the agent or (b) it falls short of every such reasonable alternative only by expected costs borne by the agent.<sup>10</sup>

**Like** MEC, COST both allows us to draw on a rich decision-theoretic tradition and promises to help unify rational and moral decision-making under uncertainty. Unlike MEC, it can both accommodate and illuminate plausible judgments about moral decision-making under uncertainty. In Section VII, I consider two broad families of response to COST. The first argues that this whole approach to deontological decision-making under uncertainty is misguided; the second accepts the approach but thinks that MEC can be salvaged. I argue, in particular, that COST earns its ‘and be-

9. See Seth Lazar, “Moral Status and Moral Options” (unpublished manuscript, Australian National University, 2016), and “Self-Ownership and Options” (unpublished manuscript, Australian National University, 2016).

10. In economics, ‘marginal’ has different uses. I use it as in the concept of ‘diminishing marginal utility’: the marginal costs and benefits of an option are the costs and benefits relative to an alternative option. This contrasts with the absolute costs and benefits, which are the option’s overall costs and benefits.

cause' clause: it doesn't just tell us which acts are subjectively permissible, it tells us *why* they are permissible.

## II. MAXIMIZING EXPECTED CHOICEWORTHINESS

Orthodox decision theory tells us (very roughly) that, when faced with practical uncertainty, we should identify our options and the possible outcomes to which each might lead, assign utilities to those outcomes and probabilities conditional on taking that option, then multiply the two numbers together before choosing the option for which the sum of those products is greatest.<sup>11</sup> This suggests the following criterion of subjective permissibility:

MEU: An act is subjectively permissible if and only if, and because, it maximizes expected utility.

Deontologists may think MEU is unavailable to them.<sup>12</sup> This would be unfortunate. Decision theorists have a rich tradition of thinking about uncertainty; deontologists do not.<sup>13</sup> One might hope that the novices could learn something from the initiates. Moreover, *prima facie*, rational and moral decision-making under uncertainty should share the same structure—after all, we often face decisions under uncertainty in which moral and prudential considerations conflict, and one would think that we could address them within a single theoretical framework.

Luckily for deontologists, their initial prejudices can be overcome. With one exception (the main subject of this article), each of deontological ethics' central commitments can be accommodated within this framework.

The first step is to recognize that 'utility' is a capacious concept, with a much broader meaning than 'individual well-being'. Unfortunately, the word carries too much baggage, so let us think in terms of 'choiceworthiness' instead and call the principle Maximize Expected Choiceworthiness.<sup>14</sup>

11. Decision theory is no more settled than ethics. This is a simplified statement of the orthodox view, and it punts on difficult questions like whether evidential or causal decision theory is correct. For two overviews, see Lara Buchak, "Decision Theory," in *Oxford Handbook of the Philosophy of Probability*, ed. Alan Hájek and Christopher Hitchcock (Oxford: Oxford University Press, 2016), 789–814; Rachael Briggs, "Normative Theories of Rational Choice: Expected Utility," in *Stanford Encyclopedia of Philosophy*, ed. Edward Zalta (Stanford, CA: Stanford University, 2014), <https://plato.stanford.edu/entries/rationality-normative-utility/>.

12. Ron Abodi, Adi Borer, and David Enoch, "Deontology, Individualism, and Uncertainty," *Journal of Philosophy* 105 (2008): 259–72; Isaacs, "Duty and Knowledge."

13. A substantial part of that decision-theoretic tradition, of course, focuses on problematizing MEU.

14. It may jar to think of an *outcome* as choiceworthy, rather than an *option*. The idea is that one outcome is more choiceworthy than another just in case you have more reason to bring it about.

MEC: An act is subjectively permissible if and only if, and because, it maximizes expected choiceworthiness.

Throughout this article, ‘choiceworthiness’ refers exclusively to moral choiceworthiness. And the ‘choiceworthiness function’ is always objective: it is a representation of our objective moral theory’s verdicts on the possible outcomes that our options might realize, ranking them on an interval scale.<sup>15</sup> According to MEC, we are subjectively permitted only to choose options with the greatest expected choiceworthiness—the highest probability-weighted average of the objective rankings of the possible outcomes.

The obvious first step is to note that an outcome can include everything that results given that I  $\phi$ . Apart from the causal consequences of  $\phi$ -ing, this also includes the fact that I  $\phi$ -d, as well as historical and relational properties such as how the outcome was brought about, what the alternative options were when I acted, and so on. This means that we can include any and all facts that a deontologist might think morally relevant to evaluating the action that brings the outcome about.<sup>16</sup>

To derive a choiceworthiness function for outcomes understood in this way, we can take one of two approaches: infer it either directly from deontic statuses or derivatively from reasons. The first approach takes the deontic verdicts of our objective theory and represents all ‘objectively permissible outcomes’ (i.e., outcomes in which I have acted objectively permissibly) as equally choiceworthy, and more so than all in which I have acted objectively wrongly. We might then rank objectively wrongful outcomes according to how seriously wrongful they are, or we could (much less plausibly) have a two-place ranking. Alternatively, we can identify the objective moral reasons for or against those outcomes and judge that one outcome is more choiceworthy than another just in case there is more moral reason in its favor.

Regardless of which of these two options we take, MEC requires that we represent all outcomes in which one has acted objectively permissibly as being equally choiceworthy.<sup>17</sup> The simplest way to see this is to note

15. This is called ‘consequentializing’ a moral theory: Amartya Sen, “Rights and Agency,” *Philosophy and Public Affairs* 11 (1982): 3–39; Broome, *Weighing Goods*; James Dreier, “Structures of Normative Theories,” *Monist* 76 (1993): 22–40; Louise, “Relativity”; Douglas W. Portmore, *Commonsense Consequentialism: Wherein Morality Meets Rationality* (Oxford: Oxford University Press, 2011).

16. It also means that the outcome includes everything that would have happened, whether or not I  $\phi$ -d. Many theories will want to disregard some of those consequences, although insofar as they will be present in the outcomes of both  $\phi$ -ing and not  $\phi$ -ing, they should cancel out. Thanks to the Editor for noting this point.

17. This will obviously mean that choiceworthiness can come apart from whatever a moral theory deems valuable: two outcomes in which you have acted objectively permissibly must be represented as equally choiceworthy, even if one is in some other sense much better than the other.

that, for MEC, objective permissibility is the limit case of subjective permissibility, in which all probabilities are 1 or 0 (and are accurate). Suppose that if I  $\phi$ , outcome  $O_\phi$  will in fact come about, and if I  $\psi$ , outcome  $O_\psi$  will in fact come about. And suppose that in either eventuality, I will have acted objectively permissibly. If we do not represent  $O_\psi$  and  $O_\phi$  as being equally choiceworthy, then MEC will generate different verdicts from our objective theory in cases when the two should converge. Suppose I am certain that  $\phi$ -ing will lead to  $O_\phi$  and  $\psi$ -ing will lead to  $O_\psi$ . So I am certain that both options are objectively permissible. But if the two outcomes are not equally choiceworthy, then only one will maximize expected choiceworthiness. The other will be subjectively impermissible, even though it is objectively permissible and my information is perfect. This is absurd: the verdicts of subjective and objective permissibility should converge when the agent is acting under full information.

Since it should be trivial to formulate a choiceworthiness function by directly grounding it in deontic statuses, I will concentrate on the approach that treats reasons as basic: one outcome is objectively more choiceworthy than another, just in case there is, in light of the facts, more moral reason to bring it about. An outcome's choiceworthiness is the degree to which it is supported by moral reasons. These reasons can include anything that one thinks matters.<sup>18</sup>

This flexibility means that a choiceworthiness function can readily represent the central features of the most plausible deontological theories.<sup>19</sup> I will not try to catalog these elements here, but consider, for example, the fact that some outcome was brought about by causing harm to a person rather than allowing it to happen or by using a person as a means or in some other way violating his rights. Each of these can be represented as a moral reason against that outcome. What's more, these reasons can be agent-relative and indeed time-relative (and perhaps there are other plausible kinds of relativity too).<sup>20</sup> They are agent-relative if their strength, or

18. Brown, focusing exclusively on objective permissibility, argues that one cannot consequentialize a satisficing theory. Perhaps this is right given his complicated (and I think unnecessary) formal architecture. But setting that architecture aside, the task is easy (although uninformative). For a satisficing theory, all permissible options must be represented as equally choiceworthy (Campbell Brown, "Consequentialize This," *Ethics* 121 [2011]: 749–71).

19. Most consequentialist criticisms of deontological decision theory focus on its incompatibility with absolutist deontology (Jackson and Smith, "Absolutist"; Huemer, "Lexical"; Isaacs, "Duty and Knowledge"). But deontology is no longer synonymous with absolutism. Almost all 'practicing deontologists' recognize that rights can be overridden in extremis.

20. I am skeptical about time-relativity—if I genuinely face a dilemma between breaching one constraint now and two identical constraints later, I should minimize my own breaches.

whether they apply, depends on who the agent is; otherwise, they are **agent neutral**.<sup>21</sup> This accommodates the deontological judgment that sometimes I must prioritize my own avoidance of some act over minimizing occurrences of that act by others. It also preserves space for special obligations grounded in promises, valuable relationships, and so on.

Suppose again that  $O_\phi$  will result if I  $\phi$ , and  $O_\psi$  will result if I  $\psi$ . Saying that  $O_\phi$  is more choiceworthy than  $O_\psi$  means that, taking all the agent's agent-neutral and agent-relative reasons into account, the balance of reasons favors  $O_\phi$ . This is obviously consistent with there being (much) greater agent-neutral reason for  $O_\psi$  than for  $O_\phi$ . This allows us to model the deontological idea, cited in the introduction, that it is sometimes impermissible to maximize the good, provided we understand 'the good' as whatever we have agent-neutral reasons to bring about.<sup>22</sup>

Deontologists may bristle at the thought that their theories can be represented with a choiceworthiness function. They might worry (and consequentialists might crow) that conceding this point means that they are *really* consequentialists.<sup>23</sup> This would be a mistake. Perhaps the critics of consequentializing are right: perhaps it at best delivers extensional equivalence without explanatory power; perhaps it obscures what motivates moral agents; perhaps consequentializing is 'gimmicky'; perhaps we can also 'deontologize' any consequentialist moral theory; and perhaps the 'compelling idea' behind consequentialism—that we should promote the good—does not transfer over to consequentialized deontological theories, except in the anodyne sense that we should do what we have most reason to do.<sup>24</sup> All these objections might be right on the money, but they would not undermine my case for deontological decision theory. My interest in consequentializing is purely instrumental. It allows me to represent my moral theory in a way that makes it amenable to decision theory. It allows me to systematically show how to balance deontological risks and opportunities when acting with imperfect information. I do not have to take a stand on whether consequentializing ultimately tells us something more about our moral theory. Perhaps all those objections are right (I think that they are, although it is hard to offer an argument either way).

21. David McNaughton and Piers Rawling, "Value and Agent-Relative Reasons," *Utilitas* 7 (1995): 31–47.

22. Obviously some people (e.g., Broome, *Weighing Goods*) simply mean by 'the good' what I mean by a choiceworthiness function. One might also think that 'the good' is a subset of what we have agent-neutral reason to bring about, in which case my point would still hold.

23. For the crowing, see Louise, "Relativity"; Oddie and Milne, "Act and Value."

24. For those criticisms, see esp. Schroeder, "Teleology"; Sergio Tenenbaum, "The Perils of Earnest Consequentializing," *Philosophy and Phenomenological Research* 88 (2014): 233–40; Paul Hurley, "Consequentializing and Deontologizing: Clogging the Consequentialist Vacuum," in *Oxford Studies in Normative Ethics*, ed. Mark Timmons (Oxford: Oxford University Press, 2013), 3:123–53; Peter Vallentyne, "Gimmicky Representations of Moral Theories," *Metaphilosophy* 19 (1988): 253–63.

But perhaps they are all wrong. The consequentializers and their opponents can fight that one out. All I need is extensional equivalence at the objective level.<sup>25</sup>

In Section VII I will stick my toe a little further into the water and argue that COST (my criterion of subjective permissibility) is not only extensionally adequate but also explanatorily powerful. But even then, I need take no stand on whether an objectivist counterpart to COST would also offer more than extensional equivalence. As it happens, I think that it would not. But nothing in my argument hangs on this view, and I will not try to defend it.

I can, however, add one nudge to the consequentializing debate. Behind many objections to consequentializing is an overriding worry that the project is pointless.<sup>26</sup> In answer to the question ‘why bother?’ I have a reply: consequentializing our deontological moral theory allows us to apply it to situations with imperfect information. It makes deontological decision theory possible. Far from being the last nail in their coffin, the possibility of consequentializing objective deontological moral theories saves deontologists from a much more serious objection—that they have nothing to say about decision-making under uncertainty.

MEC is the right place for deontologists to start. It allows us to draw on both our preferred objective moral theory and a rich tradition of thinking about uncertainty. But as I show next, we cannot stop here.<sup>27</sup>

### III. MEC AND OPTIONS TO ACT SUBOPTIMALLY

Consider the following case:

*Self-Defense.* Alice can defend herself from a lethal attack only by killing Bill and using his body to shield hers. Alice is almost certain that Bill, a business competitor, ordered the attack. Alice has no dependents and no outstanding obligations.

If Bill ordered the attack, then he culpably contributed to the wrongful threat Alice now faces. Killing him is the only way to avert that threat; it is clearly proportionate to the goal of saving Alice’s life. Killing him will not (I stipulate) have any additional adverse consequences. Killing Bill is objectively permissible.

25. Thanks to a referee for pressing me to clarify here. For a similar view on the instrumental approach to consequentializing, see Colyvan, Cox, and Steele, “Modelling,” 523.

26. See esp. Schroeder, “Teleology”; Tenenbaum, “Perils”; Hurley, “Consequentializing and Deontologizing.”

27. The most prominent advocates of MEC for deontological ethics have been consequentialists such as Oddie and Milne (“Act and Value”) and the less committal Colyvan, Cox, and Steele (“Modelling”). Their views in particular cannot deliver the right verdicts in the case in Sec. III.

If Bill did not order the attack, then killing him is not only objectively wrong, it is very seriously wrong. Killing an innocent person is bad enough; killing him to use him as a means to shield your body is even worse. Perhaps it is not absolutely prohibited—if killing Bill saved many lives, it could be objectively permissible. But it is certainly ruled out where Alice saves only herself.<sup>28</sup>

On the other hand, whether or not Bill ordered the attack, Alice is objectively permitted to let the assassin kill her. This is obviously true if Bill is innocent—indeed, Alice is required to let herself be killed in that case. But even if Bill ordered the attack, Alice is not morally required to save herself. Perhaps she would be if she had dependents or outstanding obligations. But in Self-Defense she is a freewheeling monad. And even if Bill is guilty, Alice might prefer to die than to have blood on her hands.

With these ingredients alone, we can refute MEC. In brief: killing Bill risks objective wrongdoing. Alice letting herself be killed does not. So letting herself be killed must maximize expected choiceworthiness. Killing Bill is therefore subjectively impermissible and must remain so until Alice is 100 percent confident that he ordered the hit. This is a deeply implausible result.

More slowly: to identify the permissible options according to MEC, we have to first settle on an objective choiceworthiness function for the possible outcomes of Alice's options. As we have seen, for a maximizing principle we have to represent all objectively permissible options as being equally choiceworthy.

If Bill ordered the hit, then the outcome in which Alice kills him and that in which she lets herself die each involve her acting objectively permissibly. They must therefore be ranked the same in the choiceworthiness function. We can represent this by giving each outcome the same value: X.

If Bill were innocent, then Alice letting herself die would still (obviously) be objectively permissible. Its ranking in the objective choiceworthiness function, therefore, cannot plausibly be lower than that of the outcome in which Bill ordered the hit, and Alice lets herself be killed.<sup>29</sup> So its choiceworthiness must be no lower than X. Conversely, killing Bill if he is innocent is obviously objectively impermissible and indeed much

28. If Alice has good reason to believe Bill liable, then, in virtue of that fact, her killing him might not be objectively the most egregious rights violation. But it would still be very seriously wrong (see Seth Lazar, "Risky Killing and the Ethics of War," *Ethics* 126 [2015]: 91–117).

29. One might think that it is better to let herself be killed rather than kill an innocent man than it is to let herself be killed when Bill is liable. This would only strengthen my case against MEC.

worse than killing him if he ordered the hit. So this outcome must be ranked lower than X: give it the value Y, where  $Y < X$ .

Alice is certain that letting herself be killed is objectively permissible, regardless of Bill's involvement in the attack. So its expected choiceworthiness is at least X. But if Bill might be innocent, then the expected choiceworthiness of killing him must be the probability-weighted average of X and Y. Since Y is less than X, the average of the two must also be less than X. So the expected choiceworthiness of killing Bill must be lower than that of Alice letting herself die. So killing Bill is subjectively wrong: Alice letting herself be killed is subjectively required.

This is not an acceptable result. If lethal self-defense required certainty of one's target's liability to be subjectively permissible, it would in practice always be subjectively wrong. Since it is not always subjectively wrong—this is a brute judgment, but one of which I am quite confident—we have to reject MEC. Indeed, we have to reject any approach to moral decision-making under uncertainty which models all objectively permissible options as being equally choiceworthy.<sup>30</sup>

MEC can succeed as a criterion of subjective permissibility for deontological ethics only if it can satisfy two desiderata. First, it must accurately represent the deontological theory's objective moral verdicts in a choiceworthiness function.<sup>31</sup> Second, it must obviously capture, predict, and ideally explain the theory's judgments of subjective permissibility. This case shows, however, that if MEC accurately represents objective deontological morality, then it cannot successfully capture subjective deontological morality.<sup>32</sup> If it represents objectively permissible alternatives as equally choiceworthy, then it unavoidably generates a subjective requirement for Alice to sacrifice herself, unless she is certain of Bill's liability.

Moreover, the problem generalizes, in two directions: first, take any scenario in which an agent must choose between two actions, one of which,  $\phi$ , has a risk of being objectively wrong, the other of which,  $\psi$ , is sure to be objectively permissible but involves some considerable sac-

30. One such model is the rule 'minimize the expected badness of one's wrongdoing', which represents all permissible options as being equally valuable but recognizes gradations in badness of wrongdoing. Thanks to Peter Vallentyne for raising this possibility.

31. Oddie and Milne ("Act and Value") would ditch the desideratum to accurately reflect judgments of objective permissibility, instead stipulating that the expected choiceworthiness of Alice letting herself be killed is, in this case, the same as that of killing Bill. This would imply that the expected choiceworthiness of Alice letting herself be killed must change whenever the probability that Bill is liable changes. So we have to first work out which options are subjectively permissible, then represent them as maximizing expected choiceworthiness. This is hopeless as a criterion of subjectively right action.

32. Jamie Dreier raises a different problem for MEC, also using risky cases (he focuses on supererogation). His solution is to reject maximizing in favor of satisficing: "In Defense of Consequentializing," in *Oxford Studies in Normative Ethics*, ed. Mark Timmons (Oxford: Oxford University Press, 2011), 1:97–118, 110.

rifice by the agent. Unless the agent is certain that  $\phi$  is objectively permissible—unless, that is, it is not risky at all—she will be subjectively required to  $\psi$ , no matter how great the cost to herself. Imagine, for example, that the agent can choose between performing a supererogatory action, which she knows to be objectively permissible, or one that has some risk of being objectively wrong. She would be required to perform the supposedly supererogatory action.

Second, suppose that Alice's killing Bill maximizes expected choice-worthiness and letting herself be killed does not. Then it would be subjectively wrong for her to let herself be killed. But surely if she wants to sacrifice her life for Bill's sake, she is (subjectively and objectively) permitted to do so. Whenever the agent's taking on additional expected costs to herself is suboptimal, MEC says that she is subjectively prohibited from taking on those costs. But this is wrong: we have a license to self-harm in at least some such cases.

As an aside, this has important implications for the consequentializing project. Consequentializing a moral theory should mean consequentializing its verdicts on both objective and subjective permissibility. Even if we can represent any objective moral theory within a single-ranked maximizing framework, we cannot simultaneously do so while adequately representing its subjective verdicts. Realizing extensional adequacy on one level precludes doing so on the other.<sup>33</sup>

But our interest in consequentializing is purely instrumental: it helped us to get a criterion of subjective permissibility for deontological ethics. And MEC does not allow for self-favoring and self-sacrificing agent-centered options; instead it makes unreasonably costly acts subjectively morally required. Alice does not have to forbear from using lethal defensive force unless she is certain that Bill is liable. So MEC must be wrong.<sup>34</sup>

#### IV. SATISFICING IS NOT A SOLUTION

Before presenting COST, I want to consider an apparently obvious amendment to MEC. Why not reject maximizing, in favor of satisficing? We could defend:

33. This, I think, is decisive in favor of the 'dual ranking' strategy of consequentializing; e.g., Portmore, *Commonsense Consequentialism*. In general, it would be interesting to see participants in the consequentializing debate look in more detail at subjective permissibility simultaneously with objective permissibility. For more on the dual ranking approach, see n. 49.

34. Contra, among others, Oddie and Milne, "Act and Value"; Colyvan, Cox, and Steele, "Modelling."

Satisfice Expected Choiceworthiness (SEC): An act is subjectively permissible if and only if, and because, it realizes enough expected choiceworthiness.<sup>35</sup>

This could solve Alice's problem, as well as the other cases of self-favoring and self-sacrificing options. As long as killing Bill is good enough, then it is permissible, even if Alice letting herself be killed would be better.

Satisficing faces enough well-known objections that adding more might seem like piling on.<sup>36</sup> But the shift to SEC is a common response to the problems with MEC, and considering where it goes wrong helps to establish desiderata for a more successful alternative.

The first problem is obvious but nonetheless merits emphasis. Introducing a permissibility threshold to save the deontic phenomena is unmotivated and arbitrary. Why should the threshold be there, rather than somewhere else? What makes that much (expected or otherwise) choiceworthiness enough? While satisficing models the deontological idea that sometimes it is permissible to do less than the best, it does not tie it to anything explanatory.

I discuss the second problem at greater length elsewhere, but it is worth mentioning here.<sup>37</sup> SEC renders an act subjectively impermissible if it does not yield enough expected choiceworthiness. Sometimes, however, it seems subjectively permissible for one to sacrifice one's own interests, even though doing so does not realize an expectably choiceworthy outcome, just because the costs fall only on you. SEC would rule this out, in cases in which one's self-sacrifice leaves one's action below the threshold. It would do so even in cases in which one sacrifices oneself for the sake of a smaller benefit for others.

The third problem runs still deeper, and is independently interesting. SEC cannot, without further precisification, capture plausible deontic verdicts. Suppose that Bill ordered the hit on Alice; if killing him is necessary to save herself, then it is objectively permissible. But now suppose she has a third option: by barging him out of her path, she can escape down a narrow alleyway and evade the assassin. Doing so will seriously injure Bill, but he will live. Killing him is clearly objectively impermissible in this

35. This approach to rational choice owes its name and its origins to economist Herbert A. Simon, "Rational Choice and the Structure of the Environment," *Psychological Review* 63 (1956): 129–38. For a useful overview of the debate, see Byron, *Satisficing and Maximizing*.

36. Most were already recognized in Philip Pettit, "Satisficing Consequentialism," *Proceedings of the Aristotelian Society* 58 (1984): 165–76. Bill Bradley sums up the others in "Against Satisficing Consequentialism," *Utilitas* 18 (2006): 97–108.

37. Seth Lazar, "Accommodating Options" (unpublished manuscript, Australian National University, 2016).

case.<sup>38</sup> Alice can save herself without killing anyone. All harm is bad, even when inflicted on the potentially liable. If it serves no purpose, then it is wrong. So killing Bill is wrong.

Here's the problem. If Alice lacked the option to barge Bill out of her way, killing him would be objectively permissible. It would satisfy. But with that option, it does not. And if Alice could escape without harming Bill at all, then barging would not be good enough. The only way for SEC to preserve these verdicts is to adopt contextually variable thresholds, which depend on Alice's options.<sup>39</sup> If she must kill or be killed (and Bill is liable), either option exceeds the threshold. But if her options are to let herself be killed, barge Bill out of the way, or kill him, then only the first two exceed the threshold. The threshold in SEC, then, must change depending on what Alice's options are. This seriously aggravates the arbitrariness worry, since it would mean that we cannot infer from our judgments about one case what the threshold would be for another case, with different options.

What's more, if the threshold depends on what Alice's options are, then we get further complications when Alice is unsure about her options. Suppose that she has two buttons before her. Pressing one will kill Bill, the other will paralyze him from the waist down. Either will save her life. She does not know which is which. If she presses neither, she will be killed for sure. What, then, may she do?

If Alice knew which button was which, it would be objectively permissible to either paralyze Bill or sacrifice herself but impermissible to kill Bill. Any harm more serious than paralyzing Bill would fall beneath the objective satisficing threshold. Suppose we use the same threshold for Alice's risky choice. In that case pressing either button would be impermissible, since each risks killing Bill, so its expected choiceworthiness will be lower than paralyzing him for sure and thus would not cross the objective threshold. Then Alice will be subjectively required to press neither button, and so let herself be killed, even if she is just short of certain that, say, the first button will only paralyze Bill. This is far too prohibitive.

To save SEC, we would have to adopt a distinct subjectivist threshold. But on what principled grounds could we reach that threshold? One alternative is to make the threshold the probability-weighted aver-

38. For my take on the necessity constraint, see Seth Lazar, "Necessity in Self-Defense and War," *Philosophy and Public Affairs* 40 (2012): 3–44.

39. Well, it is not the only way. Another alternative would be to individuate options by the alternatives that are available to them. Then the threshold could stay the same, but the options would be valued differently depending on what the alternatives are. I think, however, that this would just shift the arbitrariness from the threshold onto the evaluation of the options and that the result would just be a notational variant, subject to the same objections, rather than a substantively different (and more promising) view. Thanks to the Editor for this point.

age of the thresholds that would apply if her only options were kill or be killed on the one hand or paralyze or be killed on the other. But this too would not work (I relegate the details to a footnote), since it would mean that pressing either button is permissible, no matter the odds.<sup>40</sup>

SEC says that there is a threshold of expected choiceworthiness and that actions falling on or above that threshold are subjectively permissible. But it gives us no principled way to work out what that threshold is. As the last case shows, we cannot infer where the threshold should be from our knowledge of the objective threshold, combined with the probabilities at stake. What's more, as the alleyway case showed, we cannot infer where the threshold should be in one case by considering our judgments about a case with different options, since we know that the threshold must change depending on what our options are.

Barring some further technical innovation, this means that we have to establish the threshold on a case-by-case basis, without any guidance from our objective moral theory. This is no more satisfying than simply representing the subjectively permissible choice as maximizing expected choiceworthiness, without demanding any consistency in our choiceworthiness function from one case to the next and without any correspondence to the objective choiceworthiness of outcomes.<sup>41</sup> We want a criterion of subjective permissibility that not only will represent the judgments of which we are already confident but will help us form judgments about hard cases and will explain why those judgments are right. SEC cannot do this.

Perhaps we could come up with some further technical fix that is beyond my (meagre) mathematical capabilities. But why should we do so? Satisficing was introduced to preserve the subjective permission to kill or be killed in Self-Defense. It neither draws on nor explains deontological ethics. Nor does it help extend our judgments from easy cases to tricky ones. Indeed, it creates more problems than it solves (especially for options to self-sacrifice). Perhaps we should come back to satisficing if all else fails. But not before then.

In fact, I think an alternative decision rule captures all the deontic phenomena that satisficing can encompass, while being well motivated by central deontological principles, avoiding arbitrary thresholds, and being relatively simple to apply. An unmotivated and complicated tech-

40. Call the choiceworthiness of killing Bill K and that of paralyzing him P (these letters represent degrees of choiceworthiness). Suppose that button 1 has a 0.7 probability of killing, 0.3 of paralyzing, while button 2 has the complement. On this approach, if Alice presses 1, the threshold would be  $0.7K + 0.3P$ ; if button 2, then it would be  $0.3K + 0.7P$ . But that's just the expected choiceworthiness of pressing each of those buttons. It would entail that pressing either satisfices, regardless of the probabilities. This is obviously absurd.

41. Oddie and Milne, "Act and Value."

nical fix to SEC cannot merit adoption if a better-motivated and simpler alternative is available.

## V. THE FOUNDATIONS OF AGENT-CENTERED OPTIONS

My aim here is to develop a standard of subjective permissibility that not only delivers the right verdicts on cases like Self-Defense but does so for the right reasons. I don't just want a technical fix. To show that COST is well motivated, I need to at least sketch the basic deontological idea that I think underpins agent-centered options. I can then show how COST implements that idea for moral decision-making under uncertainty. I concede, of course, that we can understand the foundations of moral options in different ways and will not attempt to argue against the alternatives: most of them would also provide strong support for COST.<sup>42</sup> I also concede that consequentialists will be unmoved by my argument, since they will reject the account of moral status on which it is built. I take on that challenge elsewhere.<sup>43</sup> Here I wish only to show that one plausible and recognizably deontological foundation for moral options is neatly enacted by COST.

We should start with a particular understanding of moral status. I will not adjudicate between different accounts of what grounds that status; I care instead about its implications, of which two are key. First, in virtue of our having moral status, our interests matter. Whether our lives go well or badly can determine what we and others morally ought to do. It can give us and others reasons for action. Second, our moral status means that we matter, independently from our interests. We can determine what others ought to do, and give reasons for action, independently from the reasons for action given by our interests.

Because we have status, our interests are valuable. Utilitarians and other consequentialists are right about that.<sup>44</sup> But this does not exhaust our ability to determine what we and others ought to do. We are not mere sites for the realization of value, cells or ciphers who contribute a given quantum of suffering and flourishing to the world and are oth-

42. Most notable among them are the ideas first developed by Williams in Bernard Williams and J. J. C. Smart, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1973); then in Bernard Williams, *Moral Luck: Philosophical Papers, 1973–1980* (Cambridge: Cambridge University Press, 1981); and Samuel Scheffler's arguments in *Rejection and Human Morality* (Oxford: Oxford University Press, 1992).

43. Lazar, "Moral Status and Moral Options."

44. Perhaps Frances Kamm, Warren Quinn, and Robert Nozick (see n. 46) would deny that others' well-being can ever mean that I am required to help them. But they would agree that others' well-being can give me reasons for action. Obviously I reject their skepticism about requirements to help others.

erwise morally inert. We matter necessarily, not only contingently, depending on how our interests are affected.

For a moral theory to accommodate this idea, it must provide for at least two features: constraints and options.<sup>45</sup> If I could be harmed just to provide a marginally greater benefit to someone else, then the reasons I give would be exhausted by the contribution of my well-being to the world. I would be a mere site for the realization of value. Since I am not, I must enjoy some additional protection against being harmed, over and above the reasons given by my interests alone. Hence: constraints.

Similarly, if I had to sacrifice my interests whenever doing so realizes a marginally greater good, or if I were prohibited from self-sacrifice unless it is optimal, then the sum total of my reasons would be captured by the contribution my well-being makes to the world. I would be able to determine what I and others ought to do only in virtue of the realization in my life of some quantum of well-being. I would be a mere site for the realization of value. Since I am more than that, since I matter necessarily, I have a self-favoring option to forgo sacrificing my own interests for a marginal improvement in someone else's and self-sacrificing options to damage my own interests even when doing so is overall suboptimal.

The idea that we are not mere sites for the realization of value reflects the basic Kantian belief that moral persons are ends in themselves, beings with dignity not price. And it is the core of late twentieth-century deontological ethics, especially the work of Quinn, Kamm, Rawls, and indeed Nozick.<sup>46</sup>

The 'sites' idea helps to explain not only the grounds of moral options but also what limits them. While I am not merely a site for the realization of value, I am also such a site. And so is everyone else. We have

45. Other theorists who explicitly tie together the justification for options and constraints include David Heyd (*Supererogation* [Cambridge: Cambridge University Press, 1982], 174), Shelly Kagan (*Limits*, 207, 344 [skeptically]), and Paul Hurley ("Getting Our Options Clear: A Closer Look at Agent-Centered Options," *Philosophical Studies* 78 [1995]: 163–88).

46. Warren S. Quinn, "Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing," *Philosophical Review* 89 (1989): 287–312; Frances Myrna Kamm, "Non-Consequentialism, the Person as an End-in-Itself, and the Significance of Status," *Philosophy and Public Affairs* 21 (1992): 354–89; John Rawls, *A Theory of Justice* (Oxford: Oxford University Press, 1999); Robert Nozick, *Anarchy, State and Utopia* (Oxford: Blackwell, 1974). A recent paper by Richard Yetter Chappell ("Value Receptacles," *Noûs* 49 [2015]: 322–32) nicely critiques the separateness of persons objection to utilitarianism. I discuss this in depth in Lazar, "Moral Status and Moral Options." In brief, Yetter Chappell thinks the worry comes down to treating people as though they were fungible 'value receptacles', without regard to the genuine loss that is involved in a marginal interpersonal trade-off. I think this is a mistake. The objection does not have to do with our attitudes but with our reasons. If I am required to advance or sacrifice my interests when doing so is optimal, then I am a mere site for the realization of value—my standing as an end in myself is not independently reflected in my reasons for action.

reason to advance the interests of all beings with moral status and to respect their standing as ends in themselves. Our agent-centered options must have limits. If an action benefits others at no cost to me, then I am morally required to perform it. Indeed, if the benefit to others comes at a reasonable cost to me, then I am morally required to bear it.

As for self-sacrificing options, I think that merely acting against my own interests is almost never wrong. But there are some limits on how I can undermine my own interests. It would be wrong to sell myself into slavery or to enter into a relationship grounded in my servility, for example.<sup>47</sup> What makes such acts wrong is not the damage to my interests but my fostering particular kinds of wrongful relationships. If the only thing wrong with my  $\phi$ -ing is that it goes against my interests, as in cases of regular prudential irrationality and self-harm, then I almost always have a free rein.

As beings with moral status, our interests matter. But not only our interests matter. To reflect our standing as ends in ourselves, we must have some leeway to favor our own interests when doing so is **suboptimal** and to sacrifice them suboptimally. If we lacked that leeway, then we would be mere sites for the realization of value. But our licenses to favor and to sacrifice our interests are limited: others also have moral status, so if we can help them at a reasonable cost, we are required to do so, and sometimes damaging my own interests contravenes other moral reasons and so is impermissible. I submit that this is a quintessentially deontological idea, which any deontological criterion of subjective permissibility must accommodate, but MEC cannot.<sup>48</sup>

## VI. COST

I think we can adapt MEC to accommodate self-favoring and self-sacrificing options, in a way that both captures and illuminates intuitive verdicts on cases like Self-Defense above, tying them to the idea that we are not mere sites for the realization of value. I first present my candidate principle, then explain each of its elements, and show how it applies to these cases, before considering objections in Section VII.

47. Jean Hampton, "Selflessness and the Loss of Self," *Social Philosophy and Policy* 10 (1993): 135–65; Thomas Hurka and Esther Shubert, "Permissions to Do Less than the Best: A Moving Band," in *Oxford Studies in Normative Ethics*, ed. Mark Timmons (Oxford: Oxford University Press, 2012), 2:1–27.

48. I don't claim that this is an exclusively deontological idea. Scheffler's hybrid consequentialism, e.g., would endorse it at least in part (*Rejection*). Some versions of sophisticated consequentialism, e.g., as defended by Railton, Pettit, or Portmore, might also be able to endorse COST. See Peter Railton, "Alienation, Consequentialism, and the Demands of Morality," *Philosophy and Public Affairs* 13 (1984): 134–71; Philip Pettit, *The Robust Demands of the Good* (Oxford: Oxford University Press, 2015); Portmore, *Commonsense Consequentialism*. Thanks to a referee here.

COST: An option is subjectively permissible if and only if, and because, either (a) no other option with greater expected choiceworthiness has reasonable marginal expected costs to the agent or (b) it falls short of every such reasonable alternative only by expected costs borne by the agent.<sup>49</sup>

COST has many moving parts; let us examine each in turn.

First: *option*.<sup>50</sup> To keep things simple, let me stipulate that an option is something that it is within the agent's power to do. It must be 'epistemically available' to the agent, according to the same epistemic standard as that used in our criterion of subjective permissibility.

Second: *subjectively permissible*. Remember that subjective permissibility is an umbrella term. A full defense of COST would have to choose an epistemic standard. But I don't think that choice will affect the other details of COST.

Third: *if and only if, and because*. I argue in Section VII that COST both identifies subjectively right actions and explains why they are right. I'm more confident of the first point than of the second, but COST would still be valuable even without the 'and because' clause.

Fourth: *no other option with greater expected choiceworthiness*. According to clause (a), if an option maximizes expected choiceworthiness, it is subjectively permissible. So MEC gives a sufficient condition for subjective permissibility. This seems right, provided we remember that an option's being optimal is a function of all the agent-relative and agent-neutral reasons pertaining to the outcomes that it might realize, so keeping your promises, respecting others' rights, and acting on your special obligations to your nearest and dearest, for example, will often maximize expected choiceworthiness, even if the alternatives promise greater expected agent-neutral good.

49. This principle has some parallels to ideas about objective permissibility developed in Joshua Gert, "Requiring and Justifying: Two Dimensions of Normative Strength," *Erkenntnis* 59 (2003): 5–36; and Portmore, *Commonsense Consequentialism*. It also has affinities with dual ranking consequentialism, e.g., the self/other utilitarianism defended by Theodore Sider ("Asymmetry and Self-Sacrifice," *Philosophical Studies* 70 [1993]: 117–32) and the hybrid consequentialism in Scheffler, *Rejection*. None of those authors addresses subjective permissibility in any detail (Portmore, *Commonsense Consequentialism*, 232, touches on it in passing). I consider this approach to objective permissibility in Lazar, "Accommodating Options." Executive summary: that approach is worse motivated than COST (it's more of a technical fix, less closely tied to the underlying deontological ideas than COST is), it struggles to adequately accommodate self-sacrificing options (in particular, it makes it very difficult to say that one is merely permitted, rather than required, to kill a culpable attacker in self-defense—see n. 57 below), and it is committed to thinking of options in terms of opportunity rather than production costs (see the discussion of COST\* below, which, I think, one could not represent in existing dual ranking terms).

50. Brian Hedden, "Options and the Subjective Ought," *Philosophical Studies* 158 (2012): 343–60.

Fifth: *has reasonable marginal expected costs to the agent*.<sup>51</sup> Optimal options are permissible. If they involve no additional cost to the agent, then they can be morally required. Necessarily, a higher ranked option that has zero additional expected costs for the agent has ‘reasonable marginal expected costs’. This is the easy case (although see the discussion of COST\* below). But often the highest ranked alternatives come at a cost. In those cases, suboptimal options are also permissible if the morally higher-ranked alternatives involve unreasonable marginal expected costs to the agent. What makes a given cost unreasonable? Determining this involves comparing the good done with the cost to the agent.

Costs are setbacks to the agent’s interests. Expected costs are the probability-weighted average of those setbacks. For present purposes I think COST can remain neutral between competing theories of well-being (although its implications for some might lead us to be skeptical about that conception of well-being). Some will think that our agent-centered options should focus not only on our own interests but also on those of our nearest and dearest.<sup>52</sup> COST could be amended to reflect this view (although presumably only in clause [a], since our self-sacrificing options are more narrowly focused on costs to ourselves). For my part, I think that COST should be defined in terms of cost to the agent, and that costs to their close ones should ground agent-centered options only when they are genuinely also costs to the agent. Note, though, that we also have agent-relative reasons to aid our loved ones, and these sometimes ground permissions not to maximize agent-neutral value, insofar as they make an option more choiceworthy overall. This approach entails that sometimes our agent-relative reasons to aid our loved ones can come at an unreasonable cost with respect to our own interests, such that we can make genuinely supererogatory sacrifices for those we love. This is harder to explain when we treat the agent’s loved ones’ interests as no different from her own interests.

What do I mean by *reasonable*? At root this is, I think, a brute weighing at the same theoretical level as that which determines whether one outcome is morally preferable to another. Indeed, I think that understanding how moral advances weigh against costs to the agent requires invoking a distinct dimension of normative strength—what Frances Kamm called ‘the efforts standard’ and I have elsewhere called a rea-

51. In response to an example of Peter A. Graham’s, I think it might be necessary to add in a further qualifier here, that the costs must be in an appropriate sense *unavoidable*. This caters for the unusual but important phenomenon, identified by Graham, that sometimes the expected costs of an option might be very high, but only in virtue of knowing that one will voluntarily incur those costs at a later point.

52. Thanks to a referee here.

son's 'stringency'.<sup>53</sup> The choiceworthiness ranking, by contrast, is a function of Kamm's 'precedence standard' (which I call gravity or seriousness).

To illustrate: suppose you were a unit leader engaged in a just war, tasked with killing or capturing a squad of enemy guerrillas firing at you from within a civilian building.<sup>54</sup> You can choose between calling in air strikes and clearing the building on foot. Both options are likely to have civilian casualties (we'll assume that the civilians are not liable to be killed, but the guerrillas being targeted are). Which should you choose?

First, you must ask whether the goal achieved by killing the guerrillas is worth the likely number of civilian casualties if you choose either course of action. If not, then that course is out. This involves making a gravity comparison: is the outcome in which the civilians and the guerrillas die more choiceworthy than that in which both survive?

Assuming that, as is typical, clearing the building on foot is likely to involve fewer civilian casualties but higher risks to members of your unit, the next question is: is the reduction in civilian casualties great enough for you to be required to take the additional risk of clearing the building on foot? Working this out involves making a stringency comparison: is the additional risk to you and your soldiers counterbalanced by a proportionately large additional moral benefit? For example, if you're likely to lose two more soldiers but save ten civilian lives, are you required to go in on foot? Of course, working this out will rarely be easy—balancing competing claims never is. But one virtue of COST is that it explicitly brings out these two distinct dimensions of normative strength.<sup>55</sup>

One might be dissatisfied that COST invokes another brute weighing relation. Alas, nothing can be done about that. Working out what we are permitted and required to do involves making these two kinds of comparisons. We are well practiced at both, and we cannot expect moral theory to be more determinate than this. We can, however, expose the stringency relation to as much scrutiny as has already been granted the gravity relation. Whether a cost is reasonable comes down to more than the ranking of the outcome that it realizes. If, for example, my bearing a cost averts harm for which I am responsible, then a greater cost would be reasonable than if I were not responsible for the threatened harm. My reason to save a random stranger's life might not be stringent enough

53. Frances Myrna Kamm, "Supererogation and Obligation," *Journal of Philosophy* 82 (1985): 118–38; Seth Lazar, "Associative Duties and the Ethics of Killing in War," *Journal of Practical Ethics* 1 (2013): 3–48.

54. For more detailed discussion of cases like these, see Lazar, "Necessity"; David Luban, "Risk Taking and Force Protection," in *Reading Walzer*, ed. Yitzhak Benbaji and Naomi Sussman (New York: Routledge, 2014), 230–56.

55. These are obviously closely related to Gert's two dimensions of normative strength ("Requiring and Justifying"). I think, however, that they are not quite identical.

for even my sacrificing a hand to constitute a reasonable cost, but if I am responsible for the threat to his life, then I am required to sacrifice my **own** if that will save him. Promises and other special obligations have a similar effect on the costs we are required to bear, disproportionate in relation to the choiceworthiness of the outcomes realized.

But why *marginal*? This is a decision point in COST, and I cannot here settle which way we should go. Suppose that option  $\phi$  outranks  $\psi$  but is more costly. There are no other options. According to COST, if the difference in choiceworthiness is sufficient to make the difference in cost required, then  $\psi$  is impermissible. But we could equally well think in terms of production costs, rather than opportunity costs.<sup>56</sup> Consider, then, COST\*.

COST\*: An option is subjectively permissible if and only if, and because, either (a) no other option with greater expected choiceworthiness has *reasonable expected costs* to the agent or (b) it falls short of every such reasonable alternative only by expected costs borne by the agent.

According to COST\*, if  $\phi$  outranks  $\psi$ , and  $\phi$  involves reasonable expected costs considered in isolation from  $\psi$ , then  $\psi$  is impermissible. Conversely, if  $\phi$  and  $\psi$  are equally costly but both unreasonably so when considered in isolation, then either is permissible, even if  $\phi$  outranks  $\psi$ .

The difference between these two principles is stark: one says that, for any given cost, the agent is required to realize the most choiceworthy option; the other says that an agent cannot be required to realize an option that is, considered in isolation, unreasonably costly. Suppose, for example, that you want to donate a large amount of money to your alma mater. You could, instead, give to an organization like Deworm the World. Considered in isolation, both options involve bearing costs that you cannot reasonably be required to bear. So COST\* would say that both are permissible. But the outcome realized by donating to Deworm the World would be much more choiceworthy than the alternative—many people's lives would be substantially improved, whereas if you help your university the outcome will be good but not that good. So according to COST, you may donate to your university only if the marginal cost to you of donating to Deworm the World is unreasonable. Since the monetary cost is the same, it would have to come down to the cost of failing to support an organization that you care about. COST will therefore prove much more demanding than COST\*, which does not have to make this additional appeal.

I cannot choose between COST and COST\*. COST seems more plausible, but perhaps with some more legwork, COST\* might be a for-

56. Thanks to Jon Quong and to John Thrasher for encouraging me to think more about COST\*.

midable competitor. Settling the matter here would take us too far afield. For now, although I will continue to refer to COST, I will reserve judgment as to which one is right. Both improve on MEC, and both constitute reasonable implementations of the idea that we are not mere sites for the realization of value.

Sixth: *(b) it falls short of every such reasonable alternative only by expected costs borne by the agent.* This I think is comparatively straightforward. It provides for the option to sacrifice one's own interests, when doing so is suboptimal. A 'reasonable alternative' to  $\phi$  is an option that, as per *(a)*, outranks  $\phi$  and has reasonable marginal expected costs. Why *every* reasonable alternative? Suppose  $\phi$ -ing involves sacrificing your own interests, and every other alternative, except  $\psi$ , is better than it only with respect to your own interests. Option  $\psi$ , however, improves on it in some other morally relevant way. In this case  $\phi$ -ing is impermissible. This allows for the regular kinds of self-harming and prudential irrationality referred to in Section V, while also providing scope for some self-regarding duties. Enslaving yourself, for example, would be worse than not doing so not only in virtue of the effect on your interests but also because you would be creating an objectionable kind of relationship and showing inadequate respect for your own standing as free and equal.

COST is motivated by the idea that we are not mere sites for the realization of value. We have agent-centered options to forgo maximization when *(a)* it comes at an unreasonable cost to us and *(b)* the cost involves only our own interests (precisely how much latitude we have is hard to settle). Provided COST rules out marginal interpersonal trade-offs and the moral requirement to advance one's own interests, we're in business. We can now apply COST to Alice's case of self-defense against Bill. It can both accommodate plausible intuitions about that case and help to explain and extend them.

Alice's subjective permission to let herself be killed does not depend on the probability that Bill ordered the hit. Because of clause *(b)*, she is always permitted to let herself be killed rather than kill Bill, given that the expected costs of doing so fall only on her.<sup>57</sup> Assuming that the outcome in which Bill ordered the hit and Alice kills him is morally higher ranked

57. I think that a dual ranking approach could not accommodate this insight. Suppose that Alice is certain that Bill is culpably liable. Killing him, then, is surely the morally best thing to do: better to save her innocent life than let a guilty attacker kill her. It is also all things considered best, since it is clearly best for Alice, and nothing else is at stake. So, letting herself be killed is both morally outranked and all things considered outranked, and on the dual ranking approach it must be impermissible. But we do typically think that Alice is entitled to let herself be killed if she wishes to do so. One could argue that sacrificing herself in this case actually is morally best. But that too stretches credulity: the permission to kill culpable attackers in self-defense is not a plausible instance of the option to act suboptimally.

than the one in which Bill ordered the hit and she lets him kill her, then if Alice is certain that Bill ordered the hit, killing him will maximize expected choiceworthiness and be subjectively permissible for that reason.

As the probability that Bill ordered the hit drops, at some point killing him will no longer maximize expected choiceworthiness. Still, Alice might be subjectively permitted to kill him nonetheless, because letting herself be killed would involve unreasonable marginal expected costs to her, relative to the risks involved in killing him. However, eventually the risk of killing an innocent person grows great enough that the marginal expected cost of letting herself be killed is a reasonable one for her to have to bear, and she is subjectively prohibited from using lethal defensive force and so required to let herself be killed.

More generally, COST can help us with cases in which you must choose between a supererogatory option that you know to be objectively permissible, and a less costly option that risks wrongdoing, ~~and ones~~ in which advancing your own interests is overall optimal. Obviously COST can accommodate plausible verdicts on these cases: the first kind will be catered for by (a) and will come down to whether the expected costs to the agent are reasonable given the expected benefits realized; the second is a straightforward application of (b).

COST illuminates as well as tracks deontic verdicts. Consider Self-Defense. Apart from just helping us identify which of Alice's options is subjectively permissible depending on the probabilities, COST helps distinguish between two importantly different grounds for her action to be permissible. In one scenario, killing Bill is genuinely the morally best thing to do, given her uncertainty. In another, it's not morally best, but she cannot reasonably be expected not to bring it about. This has important implications for how we respond to Alice after the fact. For example, she would be praiseworthy if her action maximized expected choiceworthiness but not if she was avoiding an unreasonable cost—although she would probably not be blameworthy then, because she has not done anything wrong.<sup>58</sup>

Another interesting implication of COST is that an option can be subjectively supererogatory when we know that it is not objectively supererogatory. Suppose, for example, that Alice encounters a child drowning in a river. She can dive in and try to save the child. Suppose that if she successfully saves the child, the cost to her will in fact be very slight. She is already in her swimming costume, so she has no expensive suit to damage; she will suffer no injuries. However, if the rescue does prove costly, it will also fail. If she saves the child, the cost will be nil. If she fails, she will die. So, if she saves the child, then doing so was not objectively supererog-

58. Some do think that you can act permissibly and yet be blameworthy: Peter A. Graham, "A Sketch of a Theory of Moral Blameworthiness," *Philosophy and Phenomenological Research* 88 (2014): 388–409.

atory, since it is not supererogatory to save a life at little cost to **yourself** but rather morally required. If she fails: well, she has just sacrificed her life for nothing. Although not objectively supererogatory, such cases will often be subjectively supererogatory according to COST, because the expected cost of trying to save the child is high. Suppose, for example, that these two are the only possible outcomes, and they are equally likely. It certainly seems supererogatory to take a 50 percent risk of dying to save a child. Of course, we could argue that if Alice dives in and saves the child at little cost to herself, her action was merely morally required, but she herself is especially praiseworthy.<sup>59</sup> But intuitively this case puts pressure on the notion that supererogation should be calculated only in terms of actual rather than expected costs.<sup>60</sup>

## VII. OBJECTIONS

I am hunting big game here. I seek a viable deontological criterion of subjective permissibility—a set of necessary and sufficient conditions for an act's being permissible in light of an agent's nonmoral uncertainty. Objections could come from any quarter. I will not try to preempt them **all** but instead will discuss two ways to be skeptical.<sup>61</sup> The first questions whether we should even try to do deontological decision theory. The second accepts the project but asserts that some alternative criterion—perhaps even MEC—can do a better job than COST.

Some will deny that moral philosophers need say anything about what one ought to do given one's uncertainty, insisting that we should work out objective permissibility, while other more empirical disciplines can pick out decision procedures appropriate to realistic circumstances of choice. I cannot, of course, refute such all-out objectivism.<sup>62</sup> But even

59. This move is often made to address intuitions about the significance of, e.g., intentions to permissibility. See T. M. Scanlon, *Moral Dimensions: Permissibility, Meaning, Blame* (Cambridge, MA: Belknap, 2008); Judith Jarvis Thomson, "Self-Defense," *Philosophy and Public Affairs* 20 (1991): 283–310.

60. Thanks to a referee for pressing here.

61. Of necessity, I focus on objections that deontologists might make. I cannot hope to defend a broadly deontological approach to ethics against consequentialist criticism; I aim instead to show that deontologists need not be derailed by uncertainty. Consequentialists often make this objection against deontological ethics. I cannot answer their other objections here. Thanks to a referee for pressing me on this.

62. I am largely persuaded by Frank Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," *Ethics* 101 (1991): 461–82; Michael J. Zimmerman, *Living with Uncertainty: The Moral Significance of Ignorance* (Cambridge: Cambridge University Press, 2008). But I can see why others are more skeptical (Judith Jarvis Thomson, *Rights, Restitution, and Risk: Essays in Moral Theory* [Cambridge, MA: Harvard University Press, 1986]; Fred Feldman, "Actual Utility, the Objection from Impracticality, and the Move to Expected Utility," *Philosophical Studies* 129 [2006]: 49–79; Peter A. Graham, "In Defense of Objectivism about Moral Obligation," *Ethics* 121 [2010]: 88–115).

all-out objectivists must presumably say something about what Alice should do in Self-Defense, given her uncertainty. Even if we are in the realm of decision procedures and ‘what a morally conscientious agent would do’,<sup>63</sup> rather than of criteria of right action, they can presumably translate the foregoing discussion into their conceptual scheme. They also need to offer some account of how to evaluate agents’ decisions under uncertainty, to provide a standard for praise and blame. COST can help with that task too. Moreover, we need some standard by which to judge our decision procedures. One, of course, is practicality—and COST might not fare so well on this score. But we need also to ask how well they apply our objective moral theory. And this cannot simply mean that we should choose the decision procedure that leads to us most frequently acting objectively permissibly, for reasons familiar from cases such as Regan’s Whiff and Poof, Jackson’s two pills, and Parfit’s Mine Shafts.<sup>64</sup> Sometimes the action most likely to be objectively right is very clearly subjectively wrong. I think that COST provides the criterion of success at which our decision procedures should aim.<sup>65</sup>

Some deontologists, however, reject all-out objectivism and yet would also reject a decision-theoretic approach to dealing with uncertainty. They might favor a piecemeal network of pro tanto principles (perhaps similar to the ‘beyond a reasonable doubt’ standard applied in criminal law).<sup>66</sup> Or they might articulate a subjectivist theory of rights, for which our moral rights are not to actual objective performances and nonperformances by others but instead are relativized to some epistemic standard.<sup>67</sup> Again, I cannot here refute these alternatives. Indeed, the terrain is so sparsely staked out that I would have to build each theory up before knocking it down. Given that none of these approaches to uncertainty (my own included) has yet been fully developed, I think we should develop each into its most successful version before deciding which succeeds. Indeed, perhaps when they are fully developed we might find some to be notational variants on others.

But some deontologists might still be unsatisfied. They might think we have sufficient reason to reject COST, without having a better theory in hand. In particular, deontological skeptics about ‘consequentializing’ may quarrel with COST’s reliance on a choiceworthiness function for

63. Graham, “In Defense of Objectivism.”

64. Donald Regan, *Utilitarianism and Cooperation* (Oxford: Oxford University Press, 1982); Jackson, “Decision-Theoretic Consequentialism”; Derek Parfit, *On What Matters* (Oxford: Oxford University Press, 2011).

65. Thanks to a referee for help here.

66. Thanks in particular to Aaron James for helping me see this possibility.

67. See, e.g., David McCarthy, “Rights, Explanation, and Risks,” *Ethics* 107 (1997): 205–25; Zimmerman, *Living with Uncertainty*; Jonathan Quong, “Rights against Harm,” *Aristotelian Society Supplementary Volume* 89 (2015): 249–66.

outcomes. Objections will follow two complementary tracks: first, some crucial feature of deontological moral theory cannot be represented in a choiceworthiness function; second, COST is not ‘properly’ deontological—even if it gets the right verdicts on particular cases, it does so for the wrong reasons. I’ll consider them in turn.

On the first: beyond my point about the implications of representing both a theory’s subjective and objective verdicts, I have little to add to the recent literature on consequentializing. Once we take options to act suboptimally out of the choiceworthiness function, as COST does, I think we can represent every judgment of a plausible deontological theory within a choiceworthiness function. All this really means is that we can identify the performance of some acts as being morally better than others. Since an act’s performance being worse than some alternative does not entail that it is impermissible, according to COST, I think matters are relatively uncomplicated. Killing is normally worse than stealing, even if both are wrong. Traveling to **Africa** with Médecins Sans Frontières and staying at home playing PlayStation might both be permissible, but the first is definitely better. We need no more than this to get COST off the ground. It is particularly important here that my goal is not to provide a deontological decision theory for every possible deontological moral theory, but only to provide one that works for the version of deontological ethics that strikes me as most plausible.

A more specific line of attack: deontological evaluations of outcomes are not mathematically well behaved enough to deliver an adequate choiceworthiness function. In particular, we might have a partial ordering, in which some outcomes are either weakly or strongly incomparable. Two outcomes are weakly incomparable when it is not the case that one is better than the other or that they are equally good, but they are what Chang calls ‘on par’ and others call ‘roughly equal’.<sup>68</sup> Two outcomes are strongly incomparable when none of these four comparative relations obtains between them.

Strong incomparability might be a problem for COST, since it would be ‘infectious’, rendering the expected choiceworthiness of one’s options undefined.<sup>69</sup> All options involving incomparable outcomes would then necessarily be permissible, because they would not be out-ranked by a more choiceworthy option. And yet, if you genuinely believe in strong incomparability, then how could you rationally choose between options with incomparable outcomes? Doing so would presuppose comparability. I am a skeptic about strong incomparability. But COST seems to deal with it well enough: if you believe in strong incomparability you

68. Ruth Chang, “The Possibility of Parity,” *Ethics* 112 (2002): 659–88.

69. William MacAskill, “The Infectiousness of Nihilism,” *Ethics* 123 (2013): 508–20.

should think that the value of options that risk strongly incomparable considerations is undefined.

Weak incomparability raises some tricky questions that I cannot address here.<sup>70</sup> But at a first pass, COST handles it adequately. If two options will lead to outcomes that are all weakly incomparable with one another, then both options will be permissible according to clause (b), since neither has greater expected choiceworthiness than the other.

One might also worry that a deontological choiceworthiness function would not be transitive.<sup>71</sup> But even if intransitivity causes problems for rational choice, these are not distinctive to deontological decision theory. We could care about transitivity if the argument for COST rested on proving a representation theorem for a deontological choiceworthiness function, showing that adherence to COST is logically connected to abiding by some apparently inescapable axioms of rational choice. But that is not my argument; indeed, even within decision theory, the approach that uses representation theorems to justify maximizing expected utility is in disfavor.<sup>72</sup> I'm an intransitivity skeptic, and for deontologists the most plausible cases of apparent intransitivity arise from thinking about options to act suboptimally, which are not, in COST, catered for by the choiceworthiness function.<sup>73</sup> But even if the deontological choiceworthiness function exhibited intransitivity, this would not be a problem for COST.

On to the second worry: what is it for a theory to be properly deontological? COST is grounded in the idea that we are not mere sites for the realization of value. It honors the deontological idea that the right is prior to the good, on every plausible interpretation. If 'the good' is well-being, then the choiceworthiness function implied by COST will sometimes rank options that (e.g.) respect rights above those that maximize well-being even taking only agent-neutral reasons into account. If 'the good' is whatever we have agent-neutral reason to pursue, then COST allows for scenarios in which one's agent-relative reasons make the agent-neutrally best thing not the most choiceworthy option. And if 'the good' is simply the choiceworthiness function entire, then the

70. See Caspar Hare, "Take the Sugar," *Analysis* 70 (2010): 237–47; Adam Bales, Daniel Cohen, and Toby Handfield, "Decision Theory for Agents with Incomplete Preferences," *Australasian Journal of Philosophy* 92 (2013): 453–70; and Robbie Williams's contribution to this symposium.

71. Larry S. Temkin, *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning* (Oxford: Oxford University Press, 2012).

72. Christopher J. G. Meacham and Jonathan Weisberg, "Representation Theorems and the Foundations of Decision Theory," *Australasian Journal of Philosophy* 89 (2010): 641–63; Kenny Easwaran, "Decision Theory without Representation Theorems," *Philosophers' Imprint* 14 (2014): 1–30.

73. Kamm, "Supererogation and Obligation."

right is prior to the good, since our standard of subjective permissibility itself provides leeway for not pursuing the most choiceworthy option.

But perhaps I am being disingenuous. COST does not say anything like ‘act so as to treat humanity, whether in yourself or in others, always also as an end, never only as a means’. Nor does it have a constructivist cast or a contractualist framework. It says nothing about underlying justification. One might complain, then, that even if COST gets the right answers to decision-making under uncertainty, it does so for the wrong reasons. In the terms of the critics of consequentializing, it might achieve extensional equivalence without explanatory power.

I think this is wrong. COST earns its ‘and because’ clause. But I need not take a stand on whether an objective counterpart of COST would do so:

COST-OBJ: An option is objectively permissible if and only if, and because, either (a) there is no other more choiceworthy option that has reasonable marginal costs to the agent or (b) it falls short of every such reasonable alternative only by costs borne by the agent.

COST-OBJ allows us to represent an objective deontological moral theory’s verdicts on cases. And it does help explain why we are permitted to choose suboptimal outcomes—because they either involve unreasonable cost or involve costs that the agent is entitled to take on himself. But one might think that the choiceworthiness function itself is **uninformative** and that on balance COST-OBJ does not warrant its ‘and because’ clause. So be it: for the purposes of this article, I can be agnostic on this point. The explanatory power of COST does not depend on the explanatory power of COST-OBJ. I think that COST does explain what makes an act subjectively permissible, because a standard of subjective permissibility should operationalize one’s objective moral theory for decision-making under uncertainty. An act that satisfies COST best balances the risks of objective wrongdoing and the opportunities of acting objectively permissibly, given the agent’s nonmoral uncertainty. The act is made subjectively permissible by the objective moral considerations that would apply, if the facts turn out to be thus and so. We reach a verdict by representing these objective considerations with a choiceworthiness function, and then applying COST, but the ultimate grounds of subjectively permissible action are the very properties that make objectively right acts right—whatever they are. For critics of consequentializing, these are the ‘unconsequentialized’ deontological considerations that we merely represent with the objective choiceworthiness function. For friends of consequentializing, these are the reasons that ground the rankings described by COST-OBJ. Whatever your stance on that debate about objective moral theory, you can agree that COST achieves explanatory power as

well as extensional equivalence, since it details the best way of balancing risks and opportunities of objectively wrong and right action, respectively, when acting with imperfect information.<sup>74</sup>

Of course, all this presupposes that a criterion of subjective permissibility must be defined in terms of one's objective theory. If you are an all-out subjectivist, or think that subjective and objective considerations are wholly independent of one another, then COST will not merit its 'and because' clause. I reject these views: I think that there are objective moral reasons and that we should try to act on them. But I cannot fight that battle here.

The project of deontological decision theory emerges from the first line of objection relatively unscathed. There may be other ways to think about uncertainty, but this one is plausible and available to deontologists. When all of the alternatives are fully developed, we can choose between them. The next line of objection accepts the project but thinks that we can improve on COST. This is uncharted territory; if I could think of an alternative that did better than COST, I would have advanced it. One might be tempted to resurrect MEC, but the discussion in Section III decisively undermined any straightforward attempt to incorporate these options to act suboptimally within that framework. You cannot both represent all outcomes in which one has acted objectively permissibly as equally choiceworthy and get sensible verdicts on what one may do under uncertainty by maximizing expected choiceworthiness. And if you retain MEC but don't represent outcomes involving objectively permissible action as equally choiceworthy, then you would have to first decide whether a given option is subjectively permissible, then derive the objective choiceworthiness function backward from that. This is obviously a pointless task.

One possibility, though, is worth mentioning. Return to Self-Defense. One might preserve some space for the agent by insisting that her probabilities and even utilities can be imprecise. She might then be permitted to maximize expected choiceworthiness according to any of the possible probability-choiceworthiness pairs, which might give her a little leeway.<sup>75</sup> Although I am sympathetic to the idea that both probabilities and utilities are at least imprecise, ~~and perhaps~~ metaphysically vague, I don't think that solves the problem raised by Self-Defense. Alice's license to die is not a result of imprecision or vagueness. We can easily design the case with precise probabilities and utilities at stake, and we would still reject MEC's verdicts. What's more, the appeal to vagueness gets the grounds of Alice's permission wrong. She has a license to let herself be killed, not because the out-

74. Thanks to a referee for pressing me to clarify my position here.

75. Thanks to Caspar Hare and Chad Lee-Stronach for (independently) raising this point.

comes are roughly equal but because she is not a mere site for the realization of value. Appealing to vagueness here is a technical fix, which not only is extensionally inadequate but gets the grounds of our options wrong.

## VIII. CONCLUSION

I have argued for COST only indirectly. I have shown how it improves on MEC and SEC, how it enacts a basic deontological idea, and how it can deliver plausible verdicts on cases in which we have options to act suboptimally. But my primary aim has been to show deontologists that we can have a decision theory too. Indeed, to a significant degree we can draw on the resources of existing decision theory. We need not reinvent the wheel, nor is this a takeover of deontological ethics by decision theory. Deontological decision theory can, of course, learn a great deal from classical decision theory. But it also adds its own crucial, defining element, in shaping the decision rule around our options to act suboptimally. This is a happy result—happier, I think, than the suggestion that moral theory can simply be brought within the ambit of regular decision theory without the latter having to make any significant adjustments.<sup>76</sup>

COST is obviously not the last word on deontological decision theory. There are many branching points within decision theory proper (e.g., evidential vs. causal) that I have not broached. And there are substantive moral questions still to answer. COST, for example, is entirely ~~risk neutral~~, when it is quite plausible that other attitudes to risk might be morally appropriate and might be properly accommodated within the decision rule, rather than the choiceworthiness function.<sup>77</sup> We might also find a familiar deontological skepticism about aggregation hard to express within the framework of COST.<sup>78</sup> And properly accounting for vagueness might necessitate further changes.<sup>79</sup> Undoubtedly, other problems will arise.

Furthermore, deontologists who find COST unappealing can pursue a number of alternative avenues: propose an alternative standard of subjective permissibility, develop an account of subjective permissibility that is not reducible to a single criterion, or argue for all-out objectivism. Of these three possibilities, only the last has received sustained attention. There is much work to do.

76. Oddie and Milne, "Act and Value"; Colyvan, Cox, and Steele, "Modelling."

77. For this claim as applied to standard decision theory, see Lara Buchak, *Risk and Rationality* (Oxford: Oxford University Press, 2013). For one exploration of this topic in the moral context, see Buchak's contribution to this symposium.

78. Seth Lazar, "Aggregation and Risk" (unpublished manuscript, Australian National University, 2016).

79. Hare, "Take the Sugar."