

MORAL DECISION GUIDES

COUNSELS OF MORALITY OR COUNSELS OF RATIONALITY?

Holly M. Smith

MORAL AGENTS, wishing to use their moral codes to guide their decisions, are often impeded by lack of information about the circumstances and consequences of their actions. Mayor Katya's moral code directs her, in cases of financial retrenchment, to reduce the city's budget in the least damaging way. But what if she is uncertain whether it would be less damaging to cut the education budget or the public transportation budget? Many moral philosophers, contemplating questions of this sort, have concluded that the best moral theories (sometimes called "dual ought" theories) should have two tiers: a top tier stating what is *objectively* right and wrong, and a lower tier consisting of one or more decision guides designed to provide advice about what is *subjectively* best for agents to do in light of their uncertainty about what is objectively best to do. Such an agent's blameworthiness, if she does what is objectively wrong, depends in part on whether she also does what she believes to be subjectively wrong. Decision guides, then, have a strong link to blameworthiness.

In *Making Morality Work*, I recently described the kind of structure that such a two-tier moral code should exhibit.¹ This structure requires a large hierarchically organized set of decision guides to accommodate the many different kinds of uncertainty. The decision guides themselves recommend (or proscribe) acts as morally choice worthy, choice mandated, or choice prohibited. The subjectively right act is, in the simplest cases, the act recommended by the most highly ranked decision guide that is usable by the agent and suitable to her objective theory. These decision guides might include such user-friendly principles as "Do what is most likely to be objectively right," "Do what will maximize expected value," or "Do what your predecessor in office did in similar circumstances." My argument for this proposal starts from what I call the "Usability Demand," which requires that any acceptable moral theory must be usable by every agent on each occasion for decision making. In *Making Morality Work*, I argue that an

1 Smith, *Making Morality Work*.

acceptable moral theory can satisfy this demand, even if it is not always *directly* usable, so long as it is *indirectly* usable by means of a suite of appropriate decision guides that together would provide any moral agent with guidance for what to do when she wants to apply her moral theory.²

The issue before us in this paper is what the nature is of these decision guides: Are they *moral* principles of a certain sort, or are they principles of *rationality*, used here in the context of moral decision making? In my book I left this question open, and aim to resolve it in this paper. I will start by examining a recent attempt to show that these decision guides are prescriptions of rationality, not of morality.

1. PETER GRAHAM'S VIEW THAT SUBJECTIVE OUGHTS ARE RATIONAL OUGHTS, NOT MORAL OUGHTS

Peter Graham is currently the most energetic proponent of the view that decision guides are generic pragmatic principles—or, as he puts it in later writing, are principles of rationality.³ Graham is a staunch advocate of objectivism, the view that, roughly speaking, a person's moral obligations depend on all the facts about her situation except the facts concerning her beliefs or evidence about her situation.⁴ Given his commitment to objectivism, he must explain why (as he puts it) a morally conscientious person ought to perform an act that she believes is objectively wrong. This occurs in the famous Dr. Jill case (presented in table 1), in which John is afflicted with a minor but not trivial skin complaint. Treatment with Drug A would completely cure John, treatment with Drug B would partially cure him, treatment with Drug C would kill him, and giving him no treatment at all (D) would leave him permanently incurable. Dr. Jill must choose which treatment to use. Unfortunately, although Jill knows B would partially cure John and not treating him (D) would leave him permanently incurable, her evidence indicates only that Drug A and Drug C would have opposite effects, and that for each of Drug A and Drug C there is a 50 percent chance that the drug would completely cure John and a 50 percent chance that the drug would kill him.⁵

2 Smith, *Making Morality Work*.

3 Graham is far from the only adherent to this view. A recent example is offered by Muñoz and Spencer, who say, "For . . . uncertain agents, the subjective 'ought' is the proper guide to action; it is the 'ought' of rationality" ("Knowledge of Objective 'Oughts,'" 77).

4 Graham, "In Defense of Objectivism about Moral Obligation," 88–89. This is only roughly correct, but will do for purposes of this paper.

5 This version of the case, including the values in table 1, is from Zimmerman, *Living with Uncertainty*, 17–20. Zimmerman follows the case description in Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," 462–63. This type of case originated earlier, in Regan, *Utilitarianism and Co-operation*, 264–65.

Virtually everyone, including Graham, agrees that Jill, if she is a conscientious person, would choose Drug B for her patient John.

Table 1. Dr. Jill

Act	Value in Situation 1 ($p = 0.5$)	Value in Situation 2 ($p = 0.5$)	Actual Value	Expected Value
A	50*	-100	50*	-25
B	40	40	40	40*
C	-100	50*	-100	-25
D	0	0	0	0

Note: Asterisks indicate the best outcome in each column.

One way to explain this—my preferred way—is to say that while Drug A would be *objectively* right, Drug B would be *subjectively* right, and Jill ought subjectively to treat John with Drug B. Graham is highly allergic to the idea that there can be dual competing oughts—the objective and the subjective ought—partly because he believes that a decision maker needs an unequivocal answer to the questions of what she should do. It is no help to tell her that she ought to do A but in another sense she ought to do B instead.⁶ Graham argues instead that objective wrongs can be more or less serious, that killing John is a much more serious wrong than merely partially curing him, and that objectivism recommends to a conscientious agent like Jill that she not risk the more serious wrongs that Drugs A and C might incur, but instead prescribe the objectively wrong but less risky Drug B.

Despite Graham's argument, objectivism itself actually has nothing to say about what agents ought to do about risk in situations of uncertainty. It confines itself to prescriptions based on the actual facts, not on the agent's credences, probability estimates, or evidence concerning those actual facts. So how does Graham conclude that objectivism prescribes the less risky Drug B to Dr. Jill? He answers this question by arguing that in saying "Jill ought to use Drug B" we are not using the "ought" of moral obligation, but rather a different kind of "ought"—the *pragmatic* ought that is associated with ends and means.⁷

Graham holds that there are two such pragmatic "oughts"—a subjective one and an objective one. Clarifying this, he says:

According to the objective pragmatic "ought" (ought_{pragmatic (objective)}), a person ought to do something just in case doing it will bring about the

6 Graham, "In Defense of Objectivism about Moral Obligation," 95.

7 Graham, "In Defense of Objectivism about Moral Obligation," 103.

outcome, among the various outcomes from which she is choosing, she most prefers relative to her goals in acting... According to the subjective pragmatic “ought” (ought_{pragmatic (subjective)}), a person ought to do something just in case, roughly, doing so is the output that results from inputting into ... the correct decision theory ... the agent’s preference, and subjective probability, functions.⁸

So Graham concludes that if we say Jill ought to prescribe Drug *B*, we are not asserting anything about what Jill ought *morally* to have done, but instead are expressing a view about what she *pragmatically* (subjectively) ought to do.⁹ It is the very same ought, he says, that the devil might employ in saying to himself “I ought to cause a plague instead of an earthquake because that will cause more pain and suffering.”¹⁰ With a little imagination we can derive from this view a claim that all decision guides are really principles prescribing what it is pragmatically—or rationally—subjectively obligatory for an agent to do. They are not moral principles, but instead more general normative principles available to be used in connection with any type of decision, moral or otherwise, in which the agent attempts to decide what to do in light of her personal goals and epistemic limitations.

In the context of Graham’s discussion this seems to be an unhappy proposal. He has rejected the dual-oughts view about moral obligation, claiming that there cannot be both an objective and a subjective moral ought, since no agent can make a decision if faced with two conflicting types of “oughts.”¹¹ But why then should we accept the dual-oughts view about *pragmatic* oughts? On this view agents will still be faced with two conflicting types of pragmatic oughts (the pragmatic objective ought and the pragmatic subjective ought), and will be just as much at sea as an agent faced with two conflicting moral oughts.¹² Even

8 Graham, “In Defense of Objectivism about Moral Obligation,” 103.

9 Graham actually sets this up as something Jill says about herself, but for brevity I have put the words in our mouths.

10 Graham, “Moral Conscientiousness and the Subjectivism/Objectivism Debate about Moral Wrongness,” 28.

11 See also Graham, “Avoidable Harm,” 191n31.

12 Perhaps this could be resolved by Graham’s switching to a single subjective pragmatic theory that consists of one principle, “One subjectively ought to maximize expected value,” that generates recommendations both when the agent faces uncertainty and when the agent believes there is a 1.0 chance that a certain act would maximize value. But if Graham were willing to take this tack for pragmatic theories, why not take it for moral theories? Moreover, such a theory runs into another problem, which is that it is hardly clear that this theory would provide guidance for every agent, since relatively few are in a position to estimate which act would maximize expected value. Moral theories need to be supplemented by many decision guides, not just one.

worse, agents like Dr. Jill will be faced with yet a further conflict between the *pragmatic* subjective ought and the *moral* objective ought.¹³ By Graham's own lights, this should not qualify as an acceptable resolution to the agent's problem about what it is best to do.

There is another problem as well. Graham considers a case in which Jill has a momentary lapse in moral conscientiousness and prescribes Drug A. Before learning the outcome, she regains her conscientiousness, and says to herself, "I ought to have prescribed Drug B instead."¹⁴ Graham maintains that this is a pragmatic ought, not a moral ought. But he envisions an objector protesting to his appeal to pragmatic oughts as follows: "The 'ought' in Jill's thought about [what she ought to have done] is *not* a *pragmatic* 'ought.' It's clearly a moral 'ought.' It has a distinctly moral cast to it."¹⁵ Graham responds as follows: "The 'ought' in Jill's thought indeed has a moral cast to it. But that is certainly consistent with its being a pragmatic 'ought.' If the pragmatic 'ought' is, as I have indicated, an 'ought' relative to the goals of the agent in question, then it is natural that that 'ought' take on the cast of the goals to which it is relativized. The moral flavor of the 'ought' in Jill's thought, then, is easily explained by the fact that the goals Jill has ... are the ... thoroughly moral goals of the morally conscientious person."¹⁶

There is a major difficulty with this response. Graham initially defines his "pragmatic ought" in terms of what a person most prefers *relative to her goals* in acting.¹⁷ However, as observers we can identify what a person ought morally to do—either objectively or subjectively—even if acting morally is *not* one of her goals. Even if Jill were not a conscientious agent, we would say that, in light of her uncertainty about the effects of the drugs, she subjectively ought morally to choose Drug B.¹⁸ Graham's pragmatic oughts fail to reflect the essential feature of moral oughts that they are not hostage to the agent's own preferences. His proposal thus fails to capture what we mean when we say "Jill ought to use Drug B."¹⁹

In later work Graham apparently realizes that relativization of his pragmatic ought to the agent's actual goals leads to trouble. Accordingly, he revises his view.

13 Since Dr. Jill cannot know which particular action she objectively ought to do, a more accurate statement is that she will be faced by a conflict between her pragmatic subjective ought and her moral objective prohibition.

14 Graham, "In Defense of Objectivism about Moral Obligation," 103.

15 Graham, "In Defense of Objectivism about Moral Obligation," 104.

16 Graham, "In Defense of Objectivism about Moral Obligation," 104.

17 Graham, "In Defense of Objectivism about Moral Obligation," 103.

18 Furthermore, we can identify what would be morally wisest for an agent to do, even though we, as observers, have no goal that she act morally.

19 Some of the material in the foregoing paragraphs of this section appeared originally in Smith, "The Zimmerman-Graham Debate on Objectivism versus Prospectivism."

He now says that in saying Jill ought to use Drug B, we are saying that giving Drug B is what she *rationally* ought to do *if* she had the set of goals “we think ideally she ought to have, i.e., the set of goals of the morally conscientious person.”²⁰ He continues: “Given that decision theory is a theory of what it is rational to do, then it follows that [the rational] ‘ought’ is governed by the rules of some acceptable decision theory.”

But this revised proposal is also problematic. Graham claims that when we say Jill ought to give her patient Drug B, we are saying she *rationally* ought to do so, relative to the set of goals we think she ideally ought to have. The first problem here is that this new notion of “rationality” is a *substantive* normative notion, not a pure concept of rationality. What goals are those that Jill ought to have? There is no single set of ideal goals. Instead, there are many different ideal goals. There is the ideal of perfect financial achievement, the ideal of perfect athletic performance, the ideal of perfect altruism, the ideal of perfect prudence, the ideal of a perfect balance between morality and prudence, and so forth. As we normally evaluate Jill’s case, we consider what she ought *morally* to do, not what would increase her financial worth or improve her athletic performance. If Graham is right that we have an ideal in mind, it is a moral ideal. For our statement “Jill ought to give Drug B” to accurately convey this ideal, the ideal must be overtly (or contextually) expressed as part of what we mean. This means that our statement “Jill ought to give Drug B” is not a statement of a *purely* rational ought, as Graham claims. Instead, if it incorporates an ideal, it expresses a *moral* ideal, just as when her financial advisor says that “Jill ought to rebalance her portfolio to include more bonds” he means to express a financial ideal. Claiming that we are merely expressing some all-purpose ideal through the notion of what an agent *rationally* ought to do fails to capture what we actually mean, which could only be expressed by saying something like, “Morally speaking, Jill ought to give Drug B.”

My conclusion is that Graham has not argued successfully that we can interpret moral decision guides as principles spelling out what it is pragmatically or rationally obligatory to do.

2. THE CASE FOR DECISION GUIDES AS PRINCIPLES OF MORALITY

My own theory, as I mentioned at the start, says that an acceptable moral theory

20 Graham, “Moral Conscientiousness and the Subjectivism/Objectivism Debate about Moral Wrongness,” 26. He actually states this as “given that she has the set of goals we think ideally she ought to have, i.e., the set of goals of the morally conscientious person.” This still suggests we think she has these goals, whereas he needs a hypothetical to deal with the case in which the agent has alternative goals.

must have two tiers, a top tier consisting of the principles of objective rightness, and a lower tier consisting of a set of decision guides that are designed to enable agents to make decisions by reference to their objective moral theory via applying it indirectly through one of its decision guides. The idea is that if the moral theory's principle of objective rightness says, for example, "It is obligatory to maximize value," then an agent who applies this by following a decision guide that says, "It is choice mandated to maximize *expected* value" may be applying her moral code indirectly through use of this decision guide. But what actually counts as applying one's moral code indirectly is somewhat more complicated than this suggests.

For an agent to apply her moral code indirectly through use of a decision guide, the agent must employ the guide *because* she believes that it has an appropriate relationship to that objective account of right or wrong. Consider Liz, who believes act utilitarianism to be the correct account of objective right and wrong. She believes the expected utility rule to be the highest appropriate decision guide relative to act utilitarianism that she can presently use, and hence derives a prescription, via the expected utility rule, to perform act A as subjectively right relative to act utilitarianism. Liz counts as someone who "indirectly" applies act utilitarianism in deciding what to do, in part because she appropriately connects her governing moral theory to her decision guide and thence to her choice of action.

By contrast, consider Ned, who believes some act A would have greater expected utility than any other option, and derives from this a prescription to perform act A. Given only this information we cannot conclude that Ned has indirectly applied act utilitarianism as a theory of objective moral status. Ned might believe, for example, that the expected utility rule just *is* the sole account of right and wrong—he might be following Zimmerman's "Prospectivism" rather than act utilitarianism.²¹ In this case he would correctly understand himself to be *directly* applying his moral code, not indirectly applying act utilitarianism.

Or consider Katya, the aforementioned mayor who needs to decide whether to reduce city expenditures by cutting the education or the transportation budget. She knows that her predecessor, facing a similar issue, cut the transportation budget and was nonetheless reelected. Katya derives a prescription to cut the transportation budget from a decision guide recommending doing what one's predecessor did in similar circumstances. Katya believes this decision guide is appropriate for ethical egoism. But she *also* believes it is appropriate for someone trying to do what is prudentially best. Given only this information, we cannot determine whether Katya is indirectly applying ethical egoism or prudence.

21 Zimmerman, *Living with Uncertainty*.

For her to count as indirectly applying ethical egoism, she must use this decision guide *because* she believes it is appropriate to ethical egoism.

Part of an agent's employing a guide because she believes that it has an appropriate relationship to a given objective account of right or wrong is her believing that the guide delivers the *kind* of recommendation appropriate to that account. This is trickier than it may at first appear. An appropriate guide must specify not only the right kind of value to be fostered (say, utility, or honoring rights), but also the right kind of normative recommendation. Decision guides, as I see them, recommend (or proscribe) acts as choice worthy, choice mandated, or choice prohibited. The "valence" of the decision guide is clearly important. It would obviously be a mistake to try to indirectly apply the objective principle, "One *ought* to maximize utility" by following a decision guide stating, "It is choice *prohibited* to maximize expected utility."²² And the normative type of the choice mandate (or prohibition) must be appropriate as well. Consider an artist, Raul, who is uncertain whether to color a certain portion of his painting orange or blue. He wants to apply the color that will maximize aesthetic value. Not being sure which color is aesthetically best, he tries to indirectly apply his objective rule by following a decision guide that says "It is choice mandated, financially, to apply the color that will maximize expected aesthetic value." Raul might believe that he is indirectly applying his aesthetic principle, but in fact he has gone astray, since he is using a decision guide incorporating a *financial* choice mandate in order to indirectly apply an objective principle requiring him to maximize *aesthetic* value. There are different kinds of objective "oughts"—moral, legal, prudential, epistemological, financial, aesthetic, and so on—and the different objective "oughts" call for different kinds of choice mandates and different kinds of subjective "oughts." For a comprehensive normative theory to provide legitimate guidance to its users, it must include decision guides whose choice mandates and subjective oughts are properly tied to the theory's type of objective oughts. This means that the theory's decision guides must specify whether an act is morally choice mandated, legally choice mandated, epistemologically choice mandated, prudentially choice mandated, and so forth. Only if the decision guides of a *moral* theory provide *morally* choice-mandated recommendations will they form the proper basis for an agent, in trying to apply the objective moral principle indirectly, to be able to derive a recommendation for an action as subjectively morally obligatory, which is what the agent needs.

Such considerations suggest that we should adopt something like the following definition for what it is to be able to use a moral theory indirectly in making a decision:

22 This can be subtle. A decision guide stating, "It is choice prohibited to minimize expected utility" might be appropriate.

Definition 1: Ability to indirectly use a moral principle in the core sense to decide what to do: An agent *S* who is uncertain at t_i which of the acts she could perform (in the epistemic sense) at t_j is prescribed by *P* (a principle of objective moral obligation or rightness) is nonetheless able at t_i to indirectly use *P* in the core sense to decide at t_i what to do at t_j if and only if

- a. *S* believes at t_i of some act *A* that *S* could perform *A* (in the epistemic sense) at t_j ;
- b. at t_i *S* believes of act *A* that it is prescribed as morally choice mandated or choice worthy for performance at t_j by the highest-ranked decision guide (relative to *P*) usable by her at t_i ; and
- c. if and because *S* wanted, all things considered, to use principle *P* for guidance at t_i for an act performable at t_j , then her beliefs together with this desire would lead *S* to derive a prescription for *A* as subjectively morally obligatory (or as subjectively morally right) for her relative to *P*.²³

Definition 1 defines an agent's ability to indirectly use a moral principle in making a decision. Important features of this definition are its crucial stipulations (1) that the agent believes of some act *A* that it is prescribed as *morally* choice mandated (or choice worthy) by the highest-ranked decision guide, relative to her governing moral theory, that she can use, and (2) that she would derive a prescription for the act as subjectively *morally* obligatory (or right) for her, relative to *P*. If her belief or derived prescription would be phrased in terms of some other type of normativity, such as legal or prudential choice worthiness or subjective rightness, then she has gone astray.

Our conclusion should be that a comprehensive moral theory must include decision guides that enable agents to indirectly derive prescriptions from the theory. But these decision guides must be phrased in normative terms that correlate with the normative nature of the moral theory: they must be phrased in terms of moral choice worthiness, not choice worthiness of some other type of normativity. It follows that decision guides are counsels of morality, not counsels of prudence or legality or even rationality. They qualify as *moral* principles serving as fully fledged components of a comprehensive moral theory, not mere principles of rationality.

23 This definition is a simplified version, with slight changes, of one proposed in Smith, *Making Morality Work*, 280, definition 12.1. The changes involve inserting the word "morally" at key points, a necessity I overlooked in *Making Morality Work*. Subsequent definitions in that work address the ability of an agent to use a moral principle indirectly when she is uncertain about which of several decision guides is ranked highest.

3. OTHER TYPES OF DECISION GUIDES

Of course, some (but perhaps not all) decision guides may seem obviously appropriate, not only within many *moral* theories, but also within a wide variety of other normative theories that require decision guides to assist agents facing uncertainty. Such theories would include theories about what it is prudent to do, what it is legally appropriate to do, what a code of etiquette requires, what duties of religious observance require, and so forth. For example, the decision guide recommending (roughly) that one do what is most likely to be best is a decision guide that could be appropriate within any one of these alternative normative theories. These considerations may suggest that at least some of the moral decision guides are not specifically *moral* guides, but rather have a more general normative character.

But if, as I have argued, the decision guides used within a moral theory must issue prescriptions that recommend certain conduct as “morally choice worthy,” then clearly those decision guides are *not* appropriate for use within one of these other normative spheres. Someone trying to make a prudential decision is not helped by being told what it would be *morally* choice worthy for her to do. Each normative sphere requires its own type of prescriptions: what is morally choice worthy, what is prudentially choice worthy, what is epistemically choice worthy, and so forth. Nonetheless the more formal decision guides from distinct normative domains may share a general form and a good deal of their content. The general form of such a decision guide could be, “Acts of type *X* are [fill in with type of normativity] choice worthy.” Specific versions would state, “Acts of type *X* are *morally* choice worthy,” or “Acts of type *Y* are *prudentially* choice worthy,” and so forth. These decision guides differ in the kinds of normativity they recommend, although the same “type *X*” (such as “most likely to be right”) might appear in the decision guides of many domains. Note that these specific versions often differ from each other in what things they view as being valuable when they recommend, for example, that the agent maximize expected value. But they may not. Two decision guides from very different normative spheres might recognize the same things as valuable, but nonetheless issue normatively distinct recommendations. Thus a decision guide for ethical egoism may say that it is *morally* choice mandated to maximize expected personal well-being, while a decision guide for the prudential sphere may say that it is *prudentially* choice mandated to maximize expected personal well-being. These are very different recommendations, despite their common focus on the value of personal well-being.

The various normative spheres may have comprehensive theories with similar structures, requiring decision guides as well as objective principles. And

there may well be abstract templates for the content of certain, more formal decision guides that, with suitable substitutions, will serve in many such normative spheres. We might call these templates “principles of rational decision making,” but they must be adapted appropriately within each sphere in order to issue the prescriptions that are suitable for that sphere, and that can enable a decision maker to indirectly apply the relevant objective principle for the sphere in which she is operating.

4. CONCLUSION

Our question has been whether decision guides usable for indirectly applying objective moral theories should be considered as counsels of morality or counsels of rationality. I have argued, *contra* Peter Graham, that they cannot carry out their job unless it is part of their content that they recommend actions as *morally* choice mandated or choice worthy. This clearly places them in the category of moral principles rather than more neutral principles of rationality.²⁴

Rutgers University
hsmith@philosophy.rutgers.edu

REFERENCES

- Graham, Peter A. “Avoidable Harm.” *Philosophy and Phenomenological Research* 101, no. 1 (July 2020): 175–99.
- . “In Defense of Objectivism about Moral Obligation.” *Ethics* 121, no. 1 (October 2010): 88–115.
- . “Moral Conscientiousness and the Subjectivism/Objectivism Debate about Moral Wrongness.” Presented at the Workshop in Normative Ethics, University of Arizona Center for the Philosophy of Freedom, Tucson, January 16–18, 2020.
- Jackson, Frank. “Decision-Theoretic Consequentialism and the Nearest and Dearest Objection.” *Ethics* 101, no. 3 (April 1991): 461–82.
- Muñoz, Daniel, and Jack Spencer. “Knowledge of Objective ‘Oughts’: Mono-

²⁴ I am grateful for comments on previous versions of this material by the audiences at the Center for Ethics and Public Affairs at the Murphy Institute of Tulane University, the Roots of Responsibility Workshop at University College, London, and the Moral and Political Philosophy Seminar at Helsinki University, and especially to the referees for this journal.

tonicity and the New Miners Puzzle.” *Philosophy and Phenomenological Research* 103, no. 1 (July 2021): 77–91.

Regan, Donald. *Utilitarianism and Co-operation*. Oxford: Clarendon Press, 1980.

Smith, Holly M. *Making Morality Work*. Oxford: Oxford University Press, 2018.

———. “The Zimmerman-Graham Debate on Objectivism versus Prospectivism.” *Journal of Moral Philosophy* 15, no. 4 (2018): 401–14.

Zimmerman, Michael. *Living with Uncertainty*. Cambridge: Cambridge University Press, 2008.