The "Prospective View" of Obligation
Holly M. Smith

IN AN IMPORTANT NEW WORK, *Living with Uncertainty*, Michael Zimmerman seeks to provide an account of the conditions for "overall moral obligation."[1] There has been much philosophical dispute about the correct account of this concept; Zimmerman aims to resolve this debate. He characterizes the three major contenders as follows:

> *The Objective View (first formulation):*
> An agent ought to perform an act if and only if it is the best option that he (or she) has.

> *The Subjective View:*
> An agent ought to perform an act if and only if he believes that it is the best option that he has.

> *The Prospective View:*
> An agent ought to perform an act if and only if it is the option that has the greatest expectable value for the agent.[2]

To resolve the debate among these views, Zimmerman urges that we first fix our concept of "overall moral obligation." He argues that "overall moral obligation" should be understood as the kind of moral obligation with which the "morally conscientious person" is primarily concerned. When confronting some moral choice a person may ask herself, out of conscientiousness, "What ought I to do?" In this question "ought" expresses overall moral obligation. According to Zimmerman, "conscientiousness precludes deliberately doing what one believes to be overall morally wrong," although it does not "require deliberately doing, or trying to do, only what one believes to be

---

[1] Michael Zimmerman, *Living with Uncertainty* (Cambridge: Cambridge University Press, 2008).

[2] Ibid., pp. 2, 5, 39. Zimmerman proposes more complicated versions of the Prospective View, but since their features are immaterial to the issues I will raise, I confine myself to this simpler version. In his interpretation, the "expectable value" of an act for an agent is determined in part by the agent's *epistemic* probabilities for the various possible outcomes of the act, where an epistemic probability of a proposition is the degree of belief the agent is warranted in having, given the evidence available to her (pp. 19, 34, 35). Zimmerman also understands this notion to allow for the agent's uncertainty about the relevant *values* (pp. 38-39). His version of the Prospective View is intended to cover a wide range of substantive moral theories, including both purely consequentialist and also purely or partly deontological theories (pp. 3-5).
.

overall morally right, since on occasion one may find oneself forced to act while lacking any belief about the overall moral status of one's act."[3]

With this account in hand, Zimmerman argues that the Objective View and the Subjective View are each subject to fatal problems, and that the Prospective View is the correct account of moral obligation. Zimmerman, of course, has good company in advocating the Prospective View: versions of this view have become increasingly important contenders in these debates, and have been adopted by a significant number of other theorists.[4] Zimmerman's special contribution to these debates is his innovative new argument in favor of the Prospective View.

In this paper I argue that Zimmerman's core case for the Prospective View fails. He describes a case whose unusual structure was originally introduced by Donald Regan and subsequently redescribed by Frank Jackson.[5] Zimmerman's version of the case is as follows:

> *Jill's case:*
> Jill, a physician, has a patient, John, who is suffering from a minor but not trivial skin complaint. In order to treat him, she has three drugs from which to choose: A, B and C. All the evidence at Jill's disposal indicates (in keeping with the facts) that giving John Drug B would cure him partially and giving no drug would render him permanently incurable. But her evidence indicates there is a 50% probability that Drug A would cure him completely and a similar probability that it would kill him. Her evidence indicates the same thing about Drug C. In actual fact, Drug A would cure him while Drug C would kill him.[6]

The Objective View implies that Jill morally ought to give John Drug A, since it would completely cure him. Zimmerman argues that accepting this judgment is a mistake: as a conscientious person, Jill would give John Drug B, since given her evidence it would be far too risky for Jill to give him either Drug A or Drug C, whereas giving him no drug at all would certainly leave

---

[3] Ibid., pp. 2, 35.

[4] Most advocates, however, have not adopted Zimmerman's position that the "expectable" value of an act depends in part on the agent's uncertainty about *values*. For examples, see Jonathan Bennett, *The Act Itself* (Oxford: Clarendon Press, 1995), ch. 3; Richard Brandt, *Ethical Theory* (Englewood Cliffs, NJ: Prentice-Hall, Inc., 1959), p. 365; Allan Gibbard, W*ise Choices, Apt Feelings* (Cambridge, MA; Harvard University Press, 1990), p. 42; Shelly Kagan, *The Limits of Morality* (Oxford: Clarendon Press, 1989), p. 1; Graham Oddie and Peter Menzies, "An Objectivist's Guide to Subjective Value," *Ethics* 102 (April 1992): 512-533; T.M. Scanlon, *Moral Dimensions* (Cambridge, MA; Harvard University Press, 2008), p. 47; and Brad Hooker, *Ideal Code, Real World* (Oxford: Clarendon Press, 2000), section 3.1, in which he advocates a "Prospective View" version of rule utilitarianism.

[5] Donald Regan, *Utilitarianism and Co-operation* (Oxford: Clarendon Press, 1980); and Frank Jackson, "Decision-theoretic Consequentialism and the Nearest and Dearest Objection," *Ethics* 101: 461-82. The example has attracted increasing attention. See, for example, Niko Kolodny and John MacFarlane, "Ifs and Oughts," *The Journal of Philosophy*, vol. CVII, no. 3 (March 2010), pp. 115-143.

[6] Zimmerman, pp. 17-18.

him worse off than if she gave him Drug B.[7] But on the Objective View, giving John Drug B is wrong, and guaranteed to be wrong, since it is certain that one of her other options is better than Drug B. Since conscientiousness precludes deliberately doing what one believes to be morally wrong, and since Jill's choice of Drug B is the choice of a conscientious person despite the fact that she believes Drug B is guaranteed to be wrong according to the Objective View, it cannot be the case that giving Drug B is morally wrong. Zimmerman concludes that we must reject the Objective View in favor of an account according to which Jill giving Drug B is not morally wrong, but rather what she ought to do. The Prospective View provides just such an account, since (given reasonable values for the possible outcomes as shown in Figure 1) giving Drug B has the highest expectable value for Jill. Zimmerman concludes that the Prospective View is the correct account of the primary sense of moral obligation.[8]

ACT	POSSIBLE OUTCOME	VALUE	PROBABI- LITY	EXPECTABLE VALUE	ACTUAL VALUE
A	Complete cure	50	0.5	-25	50
A	Death	-100	0.5	-25	50
B	Partial cure	40	1	40	40
C	Complete cure	50	0.5	-25	-100
C	Death	-100	0.5	-25	-100
D	Permanent illness	0	1	0	0

Figure 1[9]

Of course, one alternative view of this situation is that Jill *objectively ought* to give John Drug A, but given that she does not know Drug A will cure him rather than kill him, she *subjectively ought* to give him Drug B, since this act is best in light of the non-normative properties she believes her options to have. Zimmerman rejects this type of "dual oughts" answer because it delivers an *equivocal* answer to Jill's question of what she ought to do. She wants to know, in the unique sense of "ought" with which she as a conscientious per-

---

[7] Ibid., pp. 18-19.

[8] At this point in the text Zimmerman merely says this argument "tend[s] to confirm" the Prospective View (p. 20). His overall argument in its favor includes arguments intended to show the other views are incorrect, as well as additional arguments intended to show that the Prospective View can handle well – or at least acceptably – a variety of other cases and issues. However, it is clear that he regards Jill's case as the "lynchpin" of his argument against the Objective View and in favor of the Prospective View (p. 57).

[9] The value and probability figures are from Zimmerman, ibid., pp. 19-20. As he points out, the recommendation to give John Drug B is consistent with a wide range of plausible (or even fuzzy) values and probabilities for the possible outcomes (p. 34).

son is concerned, what she ought to do. According to him, no guidance has been offered to her unless such a unique sense has been singled out.[10] Unlike the Dual Oughts View, the Prospective View offers her *univocal* advice on what to do – give John Drug B – and this act is precisely what we would all think is most reasonable for a person in Jill's position to choose.

There are two critical problems with Zimmerman's account. The first problem is that it flies in the face of a certain kind of continuity that arises in these cases. Zimmerman wants to say that a conscientious person would never deliberately do what she believes to be overall wrong. But he does not say – in fact he denies – that a conscientious person would never *risk* doing what she believes *might* be overall wrong.[11] For example, suppose Jill's evidence indicates that giving John Drug B is prospectively best in terms of what is actually valuable. It follows from the Prospective View that Jill ought to give John Drug B. But suppose also that Jill's *meta-ethical* evidence favors the truth of the Objective View rather than the Prospective View. Thus, in giving him Drug B, she believes she runs a high risk of not fulfilling her obligation.[12] Suppose Jill's evidence indicates that there is a 0.9998 chance that the Objective View is correct. Still, as a conscientious person, she would be wisest to give John Drug B, even though she runs a 0.9998 chance that in doing so she is doing what is overall wrong, since Drug A and Drug C have vastly lower expectable values. But suppose her evidence in favor of the Objective View is even stronger – it supports a 0.99998 chance that in prescribing Drug B she is doing what is overall wrong. Still, she ought to prescribe Drug B. One could go on like this indefinitely. What possible rationale could be offered for saying that a conscientious agent *cannot* have an obligation to perform an action which the agent justifiably believes to be morally wrong, but a conscientious agent *can* have an obligation to perform an action when the agent justifiably thinks there is (merely!) a 0.99998 chance that it is morally wrong? Zimmerman's entire argument depends on drawing a line between these two cases, but it is difficult to see what theoretical reason could justify insisting on such a distinction.[13] And if we cannot justify such a distinction, there seems no reason to accept the dictum that forms the lynchpin for Zimmerman's argument.

---

[10] Zimmerman, ibid., pp. 6-7, 27.

[11] Ibid., p. 59.

[12] Zimmerman describes such a case in ibid., p. 59.

[13] This is especially pressing, given that Zimmerman himself concedes that lottery paradox problems provide difficulties for attempts to define "knowledge" in terms of justificational probabilities (pp. 36-37). Several authors have recently suggested ways of combining non-normative uncertainties with normative uncertainties of this sort (see Ted Lockhart, *Moral Uncertainty and Its Consequences* (Oxford: Oxford University Press, 2000); Jacob Ross, "Rejecting Ethical Deflationism," in *Ethics* 116: 742-68; and Andrew Sepielli, "What to Do When You Don't Know What to Do," in Russ Shafer-Landau, ed., *Oxford Studies in Metaethics*, vol. 4 (Oxford: Oxford University Press, 2009), pp. 5-28. Zimmerman's defense of the Prospective View does not allow room for such considerations.

The second and worst problem is that it is possible to show that Zimmerman's Prospective View is vulnerable to the very same threat that he invokes to reject the Objective View. Consider this. The abstract structure of Jill's case is as follows. A certain Moral Standard I (the Objective View) requires a great deal of information to apply, while Moral Standard II (the Prospective View), which is responsive to the same values as Moral Standard I, requires less information to apply. The agent does not have enough information to identify which act is obligatory according to Moral Standard I, but she does have enough information to identify which act is obligatory according to Moral Standard II. She also knows that the act that is obligatory according to Moral Standard II is wrong according to Moral Standard I. Nonetheless, we and she agree that she ought (in some sense) to do what Moral Standard II recommends in this case. If we accept Zimmerman's dictum that a conscientious agent would never deliberately perform an act she believes to be wrong, we must reject Moral Standard I, because she believes the act she ought to perform is wrong according to Moral Standard I.

The problem is that we can replicate this structure with the Prospective View serving as Moral Standard I. Zimmerman himself recognizes that it requires a fair amount of information to apply the Prospective View, and that sometimes the agent cannot discover what it requires.[14] Thus there will be occasions when agents are unable to discover what the Prospective View requires and will need guidance from still another standard that is responsive to the same values but that requires less information to apply. One often recommended example of such a standard is the Minimax View, according to which an agent ought to perform an act if and only if the option minimizes the maximum possible loss of value.

Consider the following case.

> *Harry's case:*
> Harry, a physician, has a patient, Renee, with a moderately serious neurologically based tremor. Harry has three choices: he can do nothing (Act E), or he can prescribe Treatment F or Treatment G. Harry knows the possible outcomes of these various actions. Some of these outcomes involve curing Renee's tremor, while others involve loss of the use of a limb. If Harry prescribes no treatment, Renee's tremor might spontaneously disappear, or continue as is, or she might suffer loss of the use of her right foot. Harry has no way to estimate the probabilities of the various possible outcomes of the treatments or what their possible interactions with the possible cure are. However, his senior colleague has reliably informed Harry that Act E would *not* maximize expectable value, and Harry believes this.

Harry's choices can be represented by the following matrix.

---

[14] Zimmerman, ibid., pp. 70, 71, 173. In correspondence Zimmerman states that he remains unsure whether or not it could be impossible for an agent to discover what is prospectively best. Given his definition of "prospectively best," it seems transparent to me that some agents will not have the time, information or cognitive skills necessary to ascertain what is prospectively best for them on a given occasion, even if this is to be calculated on the basis of information immediately available to them, as Zimmerman requires.

ACT	POSSIBLE OUTCOME	VALUE	PROB-ABILITY	MAXIMIZE EXPECTABLE VALUE?	WORST POSSIBLE VALUE
E	Cure	10	?	No	-30
	No cure	-10	?		
	Lose use of right foot	-20	?		
	Not lose use of right foot	0	?		
F	Cure	10	?	?	-110
	No cure	-10	?		
	Lose use of left hand	-100	?		
	Not lose use of left hand	0	?		
G	Cure	10	?	?	-160
	No cure	-10	?		
	Lose use of right hand	-150	?		
	Not lose use of right hand	0	?		

Figure 2

Because Harry has no evidence regarding the probabilities of the various possible outcomes, he has no evidence supporting any figure for the expectable value of any act, although he does have excellent evidence – namely his senior colleague's statement – that not treating Renee (Treatment E) would fail to maximize expectable value.[15]

---

[15] Once the senior colleague conveys his information to Harry, Harry acquires testimonial evidence that Treatment E fails to maximize expectable value, even though he does not acquire any evidence about what the precise expectable values of the acts are. To make the senior colleague's claim intuitively plausible, one can think of the senior colleague as having evidence that (in the case of Treatments E and F only) the probabilities for Treatments E and F are as follows:

ACT	POSSIBLE OUTCOME	VALUE	PROB-ABILITY	EXPECTABLE VALUE	MAXIMIZE EXPECTABLE VALUE?	WORST POSSIBLE VALUE
E	Cure	10	0.2	-24	No	-30
	No cure	-10	0.8			
	Lose use of right foot	-20	0.9			
	Not lose use of right foot	0	0.1			
F	Cure	10	0.9	7	?	-110
	No cure	-10	0.1			
	Lose use of left hand	-100	.01			
	Not lose use of left hand	0	0.9			

Figure 3

The expectable value of Treatment E is -24, while the expectable value of Treatment F is 7. Thus Treatment E fails to maximize expectable value, since its expected value is worse than

Thus we have a situation, quite parallel to Jill's case, in which the agent lacks sufficient information to identify the best act according to Standard I (the Prospective View), but has enough information to apply the related Standard II (the Minimax View). And it seems perfectly credible that a conscientious agent in this case ought to do what Standard II recommends – Treatment E – even though he believes Treatment E is wrong according to Standard I.[16] A conscientious agent would not risk a really terrible outcome (loss of his left or right hand), even in order to avoid doing something (Treatment E) that he knows is wrong according to Standard I. On the Prospective View, using Treatment E is wrong, and guaranteed to be wrong, since it is certain that it fails to maximize expectable value. Since, according to Zimmerman, a conscientious agent would avoid doing what he believes to be overall morally wrong, but the conscientious person would choose Treatment E, Zimmerman's logic entails that Standard I – the Prospective View – is *not* the correct account of what is overall morally wrong. The reasoning here is perfectly parallel to Zimmerman's original reasoning in favor of the Prospective View. Thus, by replicating Zimmerman's argument for rejecting the Objective View in favor of the Prospective View, we have found an equally compelling argument that the Prospective View in turn should be rejected in favor of the Minimax View. On this logic, the Minimax View emerges as the unique, correct account of primary obligation. Thus Zimmerman's Prospective View can be hoisted with its own petard – the chief positive argument favoring its adoption can be redeployed to show that it should not be adopted after all.[17]

---

that of at least one other option. For some reason Harry's senior colleague does not convey these probability estimates to Harry, and Harry is unable to obtain them from the colleague or by himself.

[16] Zimmerman denies that there is any correct answer to the person who asks, "What ought I to do when I don't know what is prospectively best?" except the simple answer, "You ought to do whatever is prospectively best, whether you know what that is or not." He defends this "admittedly unhelpful" answer by saying that the "reason for rejecting the Objective View in favor of the Prospective View is *not* to find a helpful response to the question 'What ought I to do when I don't know what is actually best?' Rather, the move was dictated by the recognition that [Jill's original case] shows quite clearly that it is not in general the case that one ought, in the sense that expresses overall moral obligation, to do what is actually best" (Zimmerman, p. 71). But this misstates the case. The *only reason* to invoke the Prospective View in Jill's case is to deal with the fact that she does not have enough evidence to determine what she ought to do according to the Objective View. The Prospective View only becomes plausible precisely because it can be used *for decision guidance* in cases where the agent cannot use the Objective View. If Jill – or agents in general – were omniscient, we would have no reason to invoke the Prospective View.

[17] Someone might argue that, on Zimmerman's account of epistemic probability (which involves what beliefs are warranted relative to the agent's evidence), Harry's options do in fact have defined expectable values: before he hears from his senior colleague, all his options have *equal* expectable values. But once he's heard from his colleague, this is no longer true. Relative to Harry's evidence at that point, Treatment E is certain not to maximize expectable values (and so is certain to be wrong according to the Prospective View), and the other options then have the same expectable values as each other.

But this is clearly a wild conclusion. The Minimax View may well give the agent the best advice about what it is wise to do when his information is too highly impoverished to apply the Prospective View, but we should not conclude that it is the primary and unique account of obligation.[18] This is especially true since one can imagine replicating this argument yet again for some fourth account of obligation that is responsive to the same values but requires even less information than the Minimax View. What the redeployment of Zimmerman's core argument shows is not that the Minimax View is the correct overall account of obligation, but rather that the argument itself is defective, and cannot be used to establish any proposed account of overall obligation.[19]

Having seen two serious problems with Zimmerman's core argument in favor of the Prospective View, we can fairly conclude that this argument fails to establish the Prospective View as the primary and unique account of obligation.[20]

Holly M. Smith
Rutgers University
Department of Philosophy
hsmith@philosophy.rutgers.edu

---

[18] Of course some theorists might advocate that some alternative account, such as the Satisficing View, or the Minimax Regret View, is better than the Minimax View for these circumstances. This is immaterial to the argument in the text; the point is that *some* normative guide is needed when the Prospective View itself cannot be used.

[19] In "Making Moral Decisions," *Nous*, vol. XXII, no. 1 (March 1988): pp. 89-108; in "Deciding How to Decide: Is There a Regress Problem?" in Michael Bacharach and Susan Hurley, eds., *Essays in the Foundations of Decision Theory* (Oxford: Basil Blackwell, Inc., 1991), pp. 194-219; and in "Subjective Rightness," *Social Philosophy and Policy* 27 (summer 2010): pp. 64-110, I argue for what Zimmerman calls the "dual oughts" view, and in particular that the problem of agents' impoverished factual information can only be solved by supplementing principles of objective rightness with *multiple* principles of subjective rightness, each tailored to a different range of epistemic situations.

[20] I am grateful to comments from Michael Zimmerman and from two anonymous referees of this journal.