

I

The moral clout of reasonable beliefs

HOLLY M. SMITH

'You go to war with the Army you have, not the Army you might wish to have'
Secretary of Defense, Donald Rumsfeld¹

We must often make moral decisions armed with information less perfect than we might wish to have: our beliefs about the circumstances or consequences of our actions are mistaken, misleading, or significantly incomplete. Fortunately, decision theorists and moral philosophers have provided us with ample advice about which decisions are wisest when our information² is faulty in one of these ways. Of course, much work remains to be done in this arena. But underneath the debates about whether a decision-maker should seek to maximize her expected value, or satisfice, or minimize her maximum loss, there lurks another debate in which contenders have taken up opposing positions but have rarely tried to present a rationale for why their position is correct and their rival's position is wrong. The issue in this more fundamental debate is the question of whether what is moral depends on what the agent *actually believes*, or on what it *would be reasonable for her to believe*. Thus one proponent of maximizing expected value might argue that the expected value of an action is a function of the agent's subjective probability assignments, while a rival might argue that the expected

¹ William Kristol, 'The Defense Secretary We Have', [washingtonpost.com](http://www.washingtonpost.com/wp-dyn/articles/A132-2004Dec14.html), Wednesday, December 15, 2004, p. A33; <http://www.washingtonpost.com/wp-dyn/articles/A132-2004Dec14.html> (retrieved 26 August 2009). The full quotation is: 'As you know, you go to war with the Army you have. They're not the Army you might want or wish to have at a later time.'

² To aid exposition I shall use 'information' to include both beliefs that are true and those that are false.

value of an action is a function of the probability assignments it would be rational for the agent to have. In this paper I will examine this question, and argue that in one broad arena what is moral depends on the actual beliefs with which the agent is equipped at the time of her decision, not the beliefs she ought to have had. I shall argue that in this respect morality is parallel to war. Many political commentators believe that Secretary Rumsfeld should have built a better army before initiating the 2003 war in Iraq. But when war comes, you fight with the resources you have, not the resources you ought to have had. And in morality, when the time for decision arrives, you should decide on the basis of the beliefs you have, not the beliefs that you ought to have had.³

I. NARROWING THE QUESTION

I began by saying that our central question is ‘whether what is moral depends on what the agent *actually believes*, or on what it *would be reasonable for her to believe*’. This formulation is ambiguous, because ‘what is moral’ includes various types of moral assessment—such as the assessment of an action’s objective obligatoriness, its subjective obligatoriness, or its blameworthiness.

For our purposes, let us understand these key notions roughly as follows. *Objective obligatoriness* (or *rightness* or *wrongness*) is the deontic status that an action has in virtue of its actual circumstances and consequences. For example, the fact that Carl’s firing the gun kills Joe makes Carl’s act objectively wrong, even if Carl does not anticipate that his shooting will kill Joe. Unfortunately, decision-makers’ inaccurate or incomplete beliefs about the circumstances and consequences of their proposed actions often prevent them from seeing accurately what is objectively right, and indeed often prevent them from using many of the standard principles of objective rightness directly in making their decisions. Such agents nonetheless need moral guidance about what to do. For this reason the concept of ‘objective moral status’ needs to be

³ Of course, in war as in morality, sometimes you can defer the day of engagement. Many commentators believed that the United States ought not to have started the 2003 war in Iraq at all, but that if the United States did initiate a war, it should have waited until it had an army better prepared for the type of combat ahead. Similarly, in morality one often can and should defer a decision until one has acquired better information.

supplemented by a second type of moral status—an action’s subjective moral status—which an action has in virtue of the agent’s beliefs about its possible consequences and circumstances. For example, suppose an action is objectively wrong if it would involve breaking the agent’s promise. Sharon has promised to pick up something from the grocery store for her elderly mother, but cannot remember exactly what she promised. In this situation, if Sharon believes that buying milk is the purchase most likely to keep her promise, it is plausible that it is subjectively obligatory for Sharon to buy milk. *Subjective obligatoriness* (or *rightness* or *wrongness*) is the deontic status an action has in virtue of the agent’s beliefs about its morally relevant circumstances and consequences, such as Sharon’s belief that buying milk would be most likely to keep her promise. The point of introducing subjective moral status is to provide decision-makers with a status that they are capable of discerning, even when they cannot discern an action’s objective moral status.⁴

My primary focus in this paper is the question of whether *subjective obligatoriness* is a function of what the agent actually believes, or a function of what it would be reasonable for the agent to believe. Since an action’s subjective obligatoriness is the deontic status an action has in virtue of the agent’s ‘beliefs’ about its morally relevant circumstances and consequences, it is natural to ask whether the relevant beliefs are the agent’s actual beliefs, or the beliefs that it would be reasonable for her to have. In Carl’s case, is it subjectively obligatory for him not to fire the gun? According to the actual beliefs view, it is subjectively obligatory if he believes doing so might endanger someone. According to the

⁴ More sophisticated notions of ‘objective rightness’ and ‘subjective rightness’ are required, although the inquiry of this paper can proceed without introducing them here. I have articulated and argued for a more nuanced definition of ‘subjective rightness’ in ‘Subjective Rightness’, *Social Philosophy and Policy*, 27, No. 2 (Summer 2010), 64–110. Note, however, that some accounts of what makes an action objectively right recognize, as right-makers, facts about the agent’s beliefs (or beliefs it would be rational for him to have). Thus, for example, at one point Michael Zimmerman offers the following formulation of his preferred moral view, which is a view about what I would classify as ‘objective rightness’: ‘An agent ought to perform an act if and only if it is the option that has the greatest expected value for an agent’, where the act’s expected value depends on the degree of belief that the agent would be warranted in holding, given the evidence available to her. (Michael J. Zimmerman, *Living with Uncertainty* (Cambridge: Cambridge University Press, 2008), pp. 33–36.) Relative to this account of what is objectively right, what is subjectively right for the agent is (roughly) what is best in light of her beliefs about which action has the greatest expected value in this sense.

reasonable beliefs view, it is subjectively obligatory if it would be reasonable for him to believe it might endanger anyone, even if he does not actually believe this.

Before turning to our main question, we need to clear a little more ground. What I will call the ‘Reasonable Beliefs View’—the view that what is subjectively obligatory depends on what it would be reasonable for the agent to believe—has been articulated in a number of different ways. For example, many writers talk about ‘what it would be reasonable for the agent to believe’. But others talk about what the agent ‘would be justified in believing’, or what ‘the evidence available to the agent indicates’. Other formulations are found as well. These divide into two broad camps: those that focus on what the agent would have believed if she had carried out her *epistemic* duties,⁵ and those that focus on what the agent would have believed if she had carried out her *moral* duties.⁶

I shall take the position that even if there are epistemic duties to acquire (or have) information on matters of strictly epistemic importance, these duties are only relevant to the question of what an agent would be *morally* subjectively obligated to do if morality itself incorporates, among its duties, a duty to achieve these purely epistemic ends. Such a morality would have to address the issue of how to balance these epistemic ends against other more central moral ends, such as the ends of promoting human welfare or avoiding the violation of rights. Clearly it will sometimes or even often be the case that an individual ought, all things considered from a moral point of view, to neglect or short-change the improvement of her epistemic state in favor of securing some more important, distinctively moral, goal. If I have several inconsistent beliefs about the current president of Armenia (but no stake in this matter), I should defer correcting this objectionable epistemic situation in order to focus my attention on rescuing a toddler who has just fallen into a flooded storm-drain. Moreover, if my effecting the rescue requires me to have information about how fast the toddler is being swept away, I ought to use a quick and dirty estimation heuristic (which will deliver

⁵ For example, Michael Zimmerman, *ibid.*, p. 34.

⁶ For example, Gideon Rosen, ‘Culpability and Ignorance’, *Proceedings of the Aristotelian Society*, 108 (2002), 63; also Alexander A. Guerrero, ‘Don’t Know, Don’t Kill: Moral Ignorance, Culpability, and Caution’, *Philosophical Studies*, 136 (October 2007), 68.

an estimate in time) as opposed to an epistemically more accurate but time-consuming Bayesian method for making the estimate.

So far I have been speaking generically of a person's obligation to 'acquire additional information'. But we should recognize that there are two main methods of acquiring information. One method involves what we would normally call 'deliberation' or 'reflection'—purely mental activities such as performing mental arithmetic, drawing inferences from one's stock of beliefs, probing one's memory, putting two and two together, weighing considerations against each other, trying out a new conceptual approach to a problem, and so forth. The second method involves what we would normally call 'investigation' or 'research'—carrying out physical activities to elicit more information, such as consulting books or online resources, carrying out experiments or surveys, measuring objects or physical phenomena, asking other individuals for information, and so forth. The line between deliberation and research can be fuzzy: if I *attend* carefully to perceivable aspects of my situation, should my attending be seen as deliberation or as research? Despite the possibly unclear boundary between the two methods of inquiry, it will be useful in the subsequent discussion to have this distinction at hand. When we do not need to attend to the distinction, I shall use the term 'investigate' to refer generically to both kinds of activity.

The focus of our inquiry is the answer to the following question:

- Is what is *subjectively obligatory* for an agent to do a function of what she actually believes, or a function of what it would be reasonable for her to believe?

I shall use 'what it would be reasonable for the agent to believe' to denote all the other ways of articulating the Reasonable Beliefs View—ways that use different terminology, such as 'what evidence available to the agent indicates'. Although there are significant differences among these variants, I believe that the considerations adduced in this paper cumulatively address the viability of each of them.

Although there are several ways in which our question could be formulated, my discussion will focus on the contrast between the following two views, which I regard as the most plausible contenders for an account of what is subjectively obligatory for an agent to do:

- Reasonable Beliefs View: What is subjectively obligatory at t_i for an agent to do at t_j depends on what the agent subjectively ought to believe at t_i .
- Actual Beliefs View: What is subjectively obligatory at t_i for an agent to do at t_j depends on what the agent actually believes at t_i .

II. ARGUMENTS IN FAVOR OF THE REASONABLE BELIEFS VIEW

One way to approach the question of whether the Reasonable Beliefs View or the Actual Beliefs View is the correct account is to examine the arguments that have been given by advocates of the Reasonable Beliefs View in order to determine whether these arguments succeed.

The literature suggests several kinds of argument in favor of the Reasonable Beliefs View. Here I shall confine myself to examining only two of these.

II.A. The appeal to intuitions

The first of these is simply the appeal to intuitions about cases. Thus a proponent of the Reasonable Beliefs View might introduce the following case.

Injection: Karen, suffering from an incapacitating migraine, visits a clinic for treatment. Two drugs are standardly used for migraine relief: Drug E and Drug F. To avoid severe side-effects, Drug E must be administered by intramuscular injection, while Drug F must be administered by an intravenous drip. Physician's assistant Brad knows the importance of correctly administering each of these drugs, which have clear labels specifying the correct method of administration, and warning of the importance of using it. However, Brad has mixed up the two drugs in his mind, and administers Drug F by intramuscular injection without checking the label. As a result, Karen suffers irreversible gangrene in her arm, which requires amputation.

Clearly, Brad's injection of Drug F is objectively wrong. The proponent of the Reasonable Beliefs View of subjective rightness will also say that Brad's act is subjectively wrong. Although at the time Brad believed he was administering the drug as directed, nonetheless in failing to check the label he failed to investigate the situation as thoroughly as he ought

to have done, so his administration of the drug is subjectively wrong. Many will agree that Brad's act of injecting Drug F is subjectively wrong, and conclude that the Reasonable Beliefs View of subjective rightness is correct, while the Actual Beliefs View is incorrect.

But this conclusion is overly hasty. It is true that Brad's act of injecting Drug F is subjectively wrong—but it is sufficient for its subjective wrongness that Brad acted *precipitately* in injecting the drug. Instead of injecting the drug at time t_1 he subjectively ought to have checked the drug's label at time t_1 instead. He could, without cost, have deferred administering the drug until he checked the label, and given his proposed action's known high stakes he subjectively ought to have done so, even though he felt sure he remembered correctly how to properly administer Drug F. Thus we cannot use this case (or others involving such precipitate action) as evidence for the correctness of the Reasonable Beliefs View, since the subjective status of the action in question can be explained as arising from the fact that the agent had a subjectively better alternative at the time of action: namely, the alternative of acquiring more information.⁷

In my experience, defenders of the Reasonable Beliefs View often offer cases to support their position that on closer inspection turn out to be cases of precipitate action, like *Injection*.⁸ These cases cannot prove what Reasonable Belief advocates want them to prove. We need cases that avoid this problem. Here is one such case, based on one described by Doug Husak,⁹ that does not involve precipitate action.

Gunman: A gunman, armed with an assault rifle, bursts into a small windowless room in which a group of six faculty is conducting a

⁷ For a discussion of the importance of setting aside precipitate actions in considering the related question of culpable ignorance and blameworthiness, see Holly Smith, 'Culpable Ignorance', *The Philosophical Review*, 92 (October 1983), 543–71.

⁸ For example, Gideon Rosen's lead case for culpability in acting from culpable ignorance is someone who walks down a crowded sidewalk with his nose in a book, and knocks you over (Rosen, 'Culpability and Ignorance', 62). He states that the person with his nose in the book gains no excuse from the fact that he did not know you were there. But the action of the book-reader is either in itself a rash or precipitate action or both: either the book-reader believed that walking with his nose in a book on a crowded sidewalk was highly likely to result in a collision with someone, or he believed that the chance of this was sufficiently high that he should stop and look right now, rather than proceeding.

⁹ Described in the Rutgers Value Theory Discussion Group, March 7, 2009.

doctoral oral examination. The gunman locks the door and tosses two handguns on the table. Announcing that one of the handguns is fully loaded while the other is unloaded, the gunman tells the dissertation director, Linda, that she has one minute to inspect the guns to ascertain which one is loaded. She must then pick up whichever of the guns she chooses, aim it at the dissertation defender, and pull the trigger. If she refuses to shoot, he, the gunman, will shoot everyone in the room. Linda, flustered by this unexpected turn of events, does not inspect the guns, and after a minute the gunman announces that her time is up and she must immediately pick up a gun and fire it at the student. Linda, who served in the Army Reserve, knows that guns of type A are slightly more likely to misfire than guns of type B, and hence less likely to kill the student. She picks up Gun A and pulls the trigger. Unfortunately Gun A is the loaded gun, it fires successfully, and the student is killed.

In this case Linda clearly violates her objective and subjective obligations to gain crucial information by inspecting the two guns at time t_1 . If she had inspected them she would have believed (correctly) that Gun B is the unloaded gun. Believing this, she would then have picked up and fired Gun B, and no one would have been injured. Having failed this obligation, however, the only relevant belief she has is that Gun A is slightly more likely to misfire and thus less likely to kill the student than Gun B. In light of this belief she chooses Gun A, with tragic results. Advocates of the Reasonable Beliefs View say that a reasonable person would have inspected the guns and would have known that Gun B is the best gun to use. For them it follows that the subjectively obligatory act for Linda at time t_2 is to pick up Gun B, so in using Gun A she does what is subjectively wrong. Advocates of the Actual Beliefs View agree that a reasonable person would have inspected the guns—but they also say that the criterion for whether or not Linda's act at t_2 is subjectively obligatory is the fact that picking Gun A is best in light of what Linda *actually* believes at t_2 , regardless of what she would have believed if she had investigated as she ought to have done.

My experience is that trying to settle our dispute by trading intuitions about cases such as *Gunman* does not get us very far, as proponents of each side are captured by strong intuitions which they are unwilling to give up merely because others disagree with them. If we are to answer

the question of whether the Actual Beliefs View or Reasonable Beliefs View is true, we need to appeal not just to intuitions but to some more principled rationale.

II.B. The analogy with epistemic justification

As I have noted, we are working in territory occupied both by epistemic norms and by moral norms. This makes it tempting to transfer lessons from the epistemic norms into the realm of moral norms. One apparently relevant lesson from epistemology concerns an individual's inference of one proposition from another proposition that the individual already believes. Suppose Randy frequents unreliable anti-Semitic websites, and comes to believe that all bank presidents are Jewish. Learning that his new neighbor Josh is a bank president, he infers that Josh is Jewish. Is his belief justified? Most epistemologists hold that prior beliefs can only transmit justifiedness to new beliefs if the prior beliefs are themselves justified. An individual who starts with only an unjustified belief as a premise cannot move from it to a new justified belief, even though he uses an impeccable form of inference. More strongly, if an individual bases his belief on unjustified premises and nothing else, then the belief is unjustified.¹⁰ Thus, Randy is not justified in believing that Josh is Jewish. The lesson is that a belief is not justified if it is formed on the basis of an inference from unjustified beliefs.

From this feature of epistemic norms one could easily infer that there is a somewhat similar strong constraint within morality: if a later action arises from a morally tainted earlier action, then the later action is itself morally tainted by the agent's initial action. In particular one might infer that if an earlier action is subjectively wrong, then any subsequent action based on the earlier action must itself be subjectively wrong. This would explain any intuition that Linda is subjectively wrong to pick Gun A, since her earlier dereliction in not inspecting the guns morally infects any later action she chooses in the absence of knowledge she ought to have had.

The idea of an earlier dereliction tainting subsequent actions may be superficially attractive, although the notion of an action's being

¹⁰ See Alvin I. Goldman, 'Internalism, Externalism, and the Architecture of Justification', *The Journal of Philosophy*, 96, No. 6 (June 2009), 321.

‘based on’ an earlier action would have to be filled out satisfactorily. However, closer inspection reveals that we should reject this idea. For one thing, in many cases it would imply that *any* subsequent action that the agent performs on the basis of the earlier action is subjectively wrong. For example, suppose Judy foolishly takes out a mortgage for a house that she cannot afford. She is subjectively wrong to do so. On the tainting view, this means that she is subjectively wrong to omit to pay her monthly mortgage bill at a later date, given that she would not have had the opportunity to omit the payment if she had not foolishly taken out the mortgage in the first place. But the tainting view also implies that Judy is subjectively wrong to *pay* her mortgage bill on time, given that she would not have done so if she had not taken out the mortgage. We reject this idea: Judy now has an obligation to pay her mortgage bill. Giving the agent no acceptable option at the later stage, just because she did something subjectively wrong at an earlier stage, hardly seems the path that morality ought to take. Such cases show that we do not believe that the moral taint of a subjectively wrongful initial act infects later dependent acts and makes them subjectively wrong as well. Our general stance is that an agent is obligated to make the best of a bad situation—even one that she herself has created.

This lesson holds generally for cases in which an agent’s earlier dereliction leads to later choices. In particular, it holds for cases in which the agent’s earlier dereliction involves failing to obtain information. Linda subjectively ought to have inspected the guns. But the fact that she failed to do so does not—at least by argument from the analogy with epistemic norms about transfer of justification—imply that her subsequent action, which is based on the earlier failure to acquire information, is subjectively wrong as well. The purported analogy between epistemic norms concerning transfer of epistemic justification and moral norms concerning transfer of subjective rightness simply does not hold up under close scrutiny.

Our project in this section has been to canvass arguments offered in favor of their view by advocates of the Reasonable Beliefs View. Neither of the two arguments in favor of the Reasonable Beliefs View we have examined offers compelling support for it. It is time, then, to turn our attention to arguments showing that whatever further arguments might

be produced on its behalf, the Reasonable Beliefs View of subjective status cannot be correct.¹¹

III. PROBLEMS FOR THE REASONABLE BELIEFS VIEW

Let us turn our attention to problems that can be raised for the Reasonable Beliefs View about subjective rightness. We have stated this view as follows:

- What is subjectively obligatory at t_i for an agent to do at t_j depends on what the agent subjectively ought to believe at t_i .

I shall interpret the phrase ‘what the agent subjectively ought to believe at t_i ’ to refer to ‘what the agent would believe if she had investigated as she subjectively ought to have done’ (where ‘investigation’ may refer to either research or deliberation, as I have characterized them). For brevity I will sometimes use the phrase ‘what it would be reasonable for the agent to believe at t_i ’ to refer to this idea.¹²

In Section II I introduced the concept of subjective rightness by noting that decision-makers need moral guidance about what to do, but often lack accurate or complete beliefs about the circumstances and consequences of their proposed actions. Their defective epistemic situation renders them unable to use many standard principles of objective rightness in making their decisions. As I argued, this implies that the concept of ‘objective rightness’ needs to be supplemented by a second type of moral status—an action’s subjective moral status—which connects the values embodied in the account of objective rightness with

¹¹ Yet another sketch of an argument is offered by Gideon Rosen in ‘Culpability and Ignorance’, *ibid.* Arguing for a somewhat different point, he states two ‘More basic principles’: (a) It is unreasonable to expect people not to do what they *blamelessly believe* they are entitled to do, and (b) it is unreasonable to subject people to sanctions (such as blame) when it would be unreasonable to expect them to have acted differently (pp. 74–5; my emphasis). But clearly principle (a) would be just as intuitive if it were stated as follows, leaving out the assumption of blameless belief: It is unreasonable to expect people not to do what they *actually believe* they are entitled to do. From our point of view, Rosen’s argument assumes as part of the premise what we might have hoped it could prove. In *Who Knew?*, George Sher argues against the ‘Reasonable Beliefs’ view from another perspective entirely. See *Who Knew?* (Oxford: Oxford University Press, 2009), chapter 5.

¹² The phrase ‘What it would be reasonable for S to believe’ can be read as referring to what it would be morally permissible for S to believe, or to what it would be morally obligatory for S to believe. Cases often involve permissible beliefs rather than obligatory ones. However, for consistency I shall try to adhere to cases that can be interpreted as involving obligatory beliefs.

prescriptions that can actually be used by cognitively impeded moral agents in making their decisions. A comprehensive moral theory, then, will include both principles of objective rightness and principles of subjective rightness. The principles of subjective rightness should be designed so that they are usable by any agent facing a decision in which her options are evaluated by the principles of objective rightness.¹³ It is too much to hope (given the agent's epistemically flawed situation) that the agent will always correctly identify which act is subjectively obligatory. However, we can require that a principle of subjective rightness be usable as what I call an *internal guide* by agents. An agent is able to use a principle as an *internal guide* for deciding what to do just in case the agent would directly derive a prescription for action from the principle if she so wished.¹⁴ The agent who does not believe of any action that it would maximize utility cannot use act utilitarianism as an internal guide to her decision. However, if she believes the principle stating that an action is subjectively obligatory relative to act utilitarianism if it is the option that would maximize expected utility, and if she believes of some act A that it has this characteristic, then she is able to use this principle of subjective rightness as an internal guide to decision-making within the broad framework of act utilitarianism.

According to the Reasonable Beliefs View that we are currently examining, what is subjectively obligatory at t_1 for an agent to do at t_2 depends on what the agent subjectively ought to believe at t_1 . I shall argue that this conception has a fatal defect, because it fails to offer decision-makers principles that are usable as internal decision-guides on every occasion for decision.

III.A. Unusable decision-guides: first pass

Consider the situation of Linda in *Gunman*. She ought to have inspected the guns when the gunman threw them on the table, but she failed to do so. Hence she must now choose which gun to use without knowing which gun she would have believed to be loaded if she

¹³ More precisely: the set of principles of subjective rightness associated with a governing principle of objective rightness must be such that for each agent facing a decision there is at least one principle of subjective rightness which the agent can use to make her decision.

¹⁴ A more detailed account of internal decision-guides is offered in Holly Smith, 'Subjective Rightness', *Social Philosophy and Policy*, 27, No. 2 (Summer 2010), 64–110.

had investigated as she subjectively ought to have done. This suggests that Linda is not in a position to derive any guidance from a recommendation that tells her that what she subjectively ought to do is the action that would be best relative to what it would be reasonable for her to believe. Precisely because she failed to investigate, she does not know what it would be reasonable for her to believe, and hence cannot identify the action that would appear best in the light of that belief. Linda's situation is hardly unusual: we often fail to investigate as we ought to have done, and are subsequently not in a position to know what we would have believed if we had so investigated. Indeed, we are often uncertain whether we have investigated as fully as we ought to have done, and so are uncertain whether our resulting beliefs are ones that it is reasonable for us to hold. Even when the question is the narrower one of what belief it would be reasonable for us to hold relative to the evidence that we have, or the evidence that is accessible to us, we can be uncertain what the content of those 'reasonable' beliefs would be. (Consider all the arithmetic truths that, unknown to us, are entailed by the mathematical truths we do believe.) Hence it appears that the Reasonable Beliefs View offers us an account of subjective rightness that frequently leaves us without any guidance about what it would be best for us to do. It fails to offer decision-makers principles that are usable as internal decision-guides on every occasion for decision, and certainly falls much shorter of this goal than the Actual Beliefs View, which only requires agents to have beliefs about the content of their *actual* beliefs, not the content of the beliefs *it would be reasonable* for them to have.

Elsewhere I argue that a sufficiently sophisticated understanding of subjective rightness supports this complaint about the Reasonable Beliefs View, but makes it far less obviously true than it appears on first blush.¹⁵ Rather than spell out this argument here, I will develop another kind of argument on which the Reasonable Beliefs View fails to deliver adequate guidance.

III.B. Unusable decision-guides: second pass

Let us begin this new argument by noting that until now I have spoken as though there is only one time t_i at which an agent ought to investigate

¹⁵ In Holly M. Smith, 'Subjective Rightness' and *Making Morality Work* (manuscript).

prior to his decision at time t_j . This is clearly false, however, since an agent often has repeated opportunities to acquire information that may be relevant to an impending moral decision.¹⁶ Consider the following case, which exhibits this structure.

Election: Daria has the opportunity to vote in the governor's race in her state on November 3. There are several occasions, prior to November 3, on which Daria can investigate the qualifications of the gubernatorial candidates. On October 31 she could listen to the public radio station's analysis of the three candidates (Aviles, Blasi, and Clark); and if she did so she would conclude that Aviles is the best candidate. On November 1 she could listen to a televised debate among the candidates; and if she did so she would conclude that Blasi is best. On November 2 she could read the local newspaper's editorial column on the election; and if she did so she would conclude that Clark is best. Daria is not very interested in politics, and she finds there is something more entertaining to do on each of these occasions. In the end she does not take advantage of any of her opportunities to inform herself about the candidates. She is now standing in the voting booth on November 3, and must decide for whom to vote. Although she did not utilize any of her opportunities to investigate the candidates, she does have some impressions about their characteristics. What are Daria's investigatory obligations in this situation? We can consider her options being arrayed as follows:

October 30	October 31	November 1	November 2	November 3
	Public radio	TV debate	Newspaper	Vote

¹⁶ It is also the case that I have misleadingly spoken as though, on a single occasion, an agent has only one acceptable method of inquiry. Clearly this is false as well, as an agent may be able to choose among several potential avenues of investigation. In some situations several of these may appear equally promising, so it would be permissible for the agent to pursue any one of them, even though she has an obligation to pursue at least one. In such a situation it appears appropriate to say that it would be reasonable for an agent to believe either P or not-P, if pursuing one acceptable type of inquiry would lead her to believe P, whereas pursuing another acceptable type of inquiry would lead her to believe not-P. However, since this is readily accommodated in a revised definition of 'reasonable belief' I shall not spell it out in the text. Even if we restrict ourselves to deliberation, there may be several permissible ways to deliberate; that is, to mentally derive a conclusion from the evidence one already possesses.

Suppose watching the debate on November 1 would be (and is believed by Daria to be) the most informative method of finding out about the candidates. Suppose also that she knows her political obligations to be modest: although she is obligated to inform herself, she is not obligated to take advantage of more than one opportunity to do so. Then we could plausibly say that in the circumstances obtaining on October 30, Daria subjectively ought to watch the TV debate on November 1, since it is the best of the available opportunities. In the circumstances obtaining on October 31, it is again the case that Daria subjectively ought to watch the TV debate on November 1, and this continues to be true up to and through the time of the debate on November 1. However, in the circumstances obtaining on November 2, after she has failed to watch the debate, Daria subjectively ought to read the newspaper.¹⁷

For whom is she subjectively obligated to vote on November 3, when she is standing in the voting booth, having failed to take advantage of any of her opportunities to inform herself? According to the Reasonable Beliefs View, this depends on what a reasonable person in her circumstances would believe about the candidates. The Reasonable Beliefs View could provide an answer to what Daria ought to do on November 3 by adopting either of two approaches to this question. The first approach takes a *maximally idealistic* view about what standard people should be held to in circumstances where they have failed to avail themselves of any investigative opportunity. According to this approach, what a reasonable person in her circumstances would believe depends on what she would have believed if she had taken advantage of her *best* opportunity to investigate.¹⁸ Let us call this the 'Maxi-Idealist' approach. Since Daria's best opportunity was to watch the TV debate, which would have led her to believe that Blasi is best, this approach would result in the conclusion that she subjectively ought to vote for

¹⁷ Moral Actualists and Possibilists might disagree on what, as of October 30, Daria ought to do, or on the grounds on which she ought to do it. This would depend on details of her psychology, which for the purposes of this paper I will leave aside.

¹⁸ My characterization of both the Maxi-Idealist and the Mini-Idealist position assumes that the Reasonable Beliefs advocate can provide an account of what 'an occasion to investigate' is, in terms of what the relevant issues to be investigated are. It is not clear that this will be an easy task, especially since any given occasion to investigate one fact relevant to the impending decision may not afford the possibility of investigating other relevant facts, and it will hardly be easy to individuate 'facts'.

Blasi, since this is the vote that would be subjectively best if she believed, as she ought to (on the Maxi-Idealist approach), that Blasi is the best candidate.¹⁹

The second approach takes a somewhat less idealistic view of what standard people should be held to in circumstances where they have failed to avail themselves of any investigative opportunity. According to this approach, what a reasonable person in Daria's circumstances would believe depends on what she would have believed if she had taken advantage of her *last possible opportunity* to investigate. Let us call this the 'Mini-Idealist' approach. Assuming that Daria's last possible opportunity to investigate was her opportunity on November 2 to read the newspaper, which would have led her to believe that Clark is best, it appears that on the Mini-Idealist approach she should vote as she would if she believed Clark to be the best candidate.²⁰

The Maxi-Idealist approach to multiple sequential opportunities to investigate

It appears to me that the Maxi-Idealist approach to this question is the one most consistent with the idealistic spirit of the Reasonable Beliefs View. Nonetheless, this approach suffers a fatal problem. To see the problem, we need to consider a case with features slightly different from those of *Election*.

In the kind of case we have been considering, whether or not an agent investigates a situation as she subjectively ought to do typically makes a difference to what she believes at the time of the decision itself. For example, if Daria listens to the radio analysis she will come to believe that Aviles is the best candidate, whereas if she watches the televised debate she will come to believe that Blasi is the best candidate.

¹⁹ For the sake of argument I am temporarily waiving questions of whether Daria would have any belief about what she would believe if she had taken advantage of her best (or her last) opportunity to investigate, as of course she might not. This issue is addressed in Section III.A. Note that we can define a supererogatory investigatory act for Daria: namely, her taking advantage of *all three* of her opportunities. However, an account of what it would be reasonable for her to believe in terms of what she would believe if she had engaged in supererogatory investigation seems beyond the idealistic hopes of even the most committed Reasonable Beliefs advocate.

²⁰ Here also I am setting aside questions of whether Daria would have any belief about what she would believe if she had taken advantage of her last opportunity to investigate.

However, reflecting on these sequential-investigation cases reminds us that an agent's investigation can affect not only her beliefs but also *the nature of the decision situation* itself. Consider the following case.

Space medicine: Alex is the new medic for the space station. On Monday, prior to taking off for the space station, Alex subjectively ought to investigate which antibiotic is most effective under zero-gravity conditions. If Alex investigated this question on Monday he would discover that antibiotic A is most effective, followed in order by antibiotics C and B. Having discovered this, he would pack antibiotics A, B, and C in his medicine kit. However, he fails to investigate this question on Monday. Instead, he relies on his background beliefs that (under normal gravity conditions) antibiotic A is very ineffective, whereas B and C are effective, with B being more effective than C. Given these beliefs, Alex packs only antibiotics B and C in his medicine kit. On Wednesday, once he and the other crew-members are at the space station, he has another opportunity to investigate the question of effective antibiotics during the once-a-day opportunity to communicate with ground control at Houston. When he asks the ground-based physicians which antibiotic is most effective in zero-gravity conditions, they reply that antibiotic A is most effective. Unfortunately, a sudden magnetic storm interrupts communication before they can address the efficacy of antibiotics B and C. Early on Thursday a virulent infection breaks out among the crew. Not having packed antibiotic A, Alex must choose between using antibiotic B or antibiotic C to treat the infection. The magnetic storm has continued, preventing all communication with Houston for an unpredictable period of time, and Alex must decide what to do without further advice from ground-based physicians.

In this case, Alex's failure to investigate on Monday makes a difference—not only to the beliefs he has, but also *to the situation he is in* on Thursday. Had he investigated on Monday, he would have believed antibiotic A to be the best for use in zero-gravity conditions. Believing that, he would have included antibiotic A in his medical kit, and he would now be able to use it for the sick crew-members. But he did not investigate, and so he does not have the option of using antibiotic A.

According to the Maxi-Idealist approach, what a reasonable person in Alex's circumstances would believe depends on what he would have

believed if he had taken advantage of his *best* opportunity to investigate. In this case, Alex's best opportunity occurred on Monday.

Suppose Alex tries to follow the following principle of subjective rightness:

- P: An act Y is subjectively obligatory if it would be reasonable for the agent to believe that Y is available and would maximize expected value.

Because of the Wednesday message from ground-based physicians, Alex knows that if he had investigated as he subjectively ought to have done on Monday, then (a) he would have learned that antibiotic A is the antibiotic most likely to be effective in space, (b) he would have included A in his medicine kit, (c) antibiotic A would have been available to him, (d) he would have believed that A is available to him, and (e) he would have believed using antibiotic A for the crew's infection to be the action that would maximize expected value. Thus using antibiotic A is prescribed by P, and Alex knows this. But of course he *cannot* use antibiotic A, and he knows he cannot use A, since A is not actually available on the space station.²¹

What this shows is that on the Maxi-Idealist approach to defining 'reasonable beliefs' in certain cases where an agent has but fails to pursue multiple opportunities to investigate, the Reasonable Beliefs View (as expressed in principle P) fails to provide Alex with a usable decision-guide, since it prescribes as subjectively obligatory an action that he believes is not available to him.²² And this will often be true of agents

²¹ A proponent of the Reasonable Beliefs View might suggest we avoid this problem by stipulating that principle P applies only to those acts that are genuinely available to the agent at the time of decision. But this would be a mistake: one of the types of information about which agents are obliged to inquire is *which options will be available to them*. I have found no acceptable way of defining the subjectively right act, within the Reasonable Beliefs framework, that both (a) recognizes that agents must investigate which acts will be available to them, and (b) does not identify as subjectively right an act that is available but which it would not be reasonable for the agent to believe is available.

²² Some commentators will claim that *Space medicine* simply shows that the Reasonable Beliefs view is wrong because it can require an agent, such as Alex, to perform an action which it is not possible for him to perform. This view is not correct. Although it may be true that 'S objectively ought to do A' entails 'S can do A', as many have believed, 'ought' does not entail 'can' for subjective obligation. (Consider the case of an agent who believes she ought to turn off the alarm clock so that it will not waken her husband, but who is unable to perform this action because, unbeknownst to her, she has suffered a stroke during the night. Nonetheless she has a subjective obligation to turn off the alarm, even though she has no such objective obligation.)

whose choice of whether or how to investigate alters their subsequent decision situations in ways that affect which acts are available to them. Let us call these *Altered Option* cases. Alex's case is not an unusual one; whether or how we make inquiries into an upcoming decision frequently affects what course of action we take before the decision arises, and may affect what we actually have the opportunity to do at the time of the contemplated original decision.²³

The Mini-Idealist approach to multiple sequential opportunities to investigate

The Maxi-Idealist approach says that what a reasonable person would believe is what she would have believed if she had taken advantage of her *best* opportunity to investigate. We have now seen that this approach must be rejected, since it entails that the Reasonable Beliefs View cannot

A plausible parallel principle would be 'S subjectively ought to do A' entails 'S *believes* she is able to do A'. It is the latter principle that is relevant to Alex's case, and as the case shows, he does not believe that he is able to use antibiotic A, so he cannot be under a subjective obligation to do so. Note that the Maxi-Idealist approach causes trouble even in examples such as *Election* in which the agent's situation does not change because she has failed to investigate. I said above in the text that it is plausible to say that if Daria fails to avail herself of her best opportunity to investigate (watching the TV debate on November 1), then, in the circumstances obtaining on 2 November—which include the fact that she has failed to watch the debate—Daria subjectively ought to read the newspaper. But on the Maxi-Idealist approach this does not seem to be correct. What Daria subjectively ought to do on November 2 depends on what a reasonable person in her position would believe on November 2. However, a reasonable person would have watched the TV debate on November 1, would believe that she had watched the debate, would have arrived at a belief regarding who is the best candidate, and would believe that there is no need for her to avail herself of the less promising avenue of inquiry available on November 2: namely, reading the newspaper. Thus it seems *false* that Daria subjectively ought to read the newspaper on November 2. In other words, the Maxi-Idealist approach implies that agents have no obligation to correct their malfeasance—to pursue subsequent opportunities to investigate once they have unreasonably missed their best opportunity to investigate. This is surely a result of the Reasonable Beliefs View with which its advocates do not want to be stuck.

²³ Some of my interlocutors have agreed that Alex cannot be required to use antibiotic A, and have responded that this very problem is what has moved them to adopt the interpretation of 'reasonable beliefs' that identifies them with 'beliefs indicated by the evidence available to the agent', or 'beliefs indicated by the evidence possessed by the agent'. The subsequent arguments in the text show that these interpretations are equally vulnerable to problems. But the very acknowledgment that a Reasonable Beliefs View advocate should adopt the 'evidence' interpretation rather than the 'what research would have shown' interpretation shows that the Reasonable Beliefs View advocate admits that usability is an important criterion for assessing accounts of subjective rightness. Since an agent can possess evidence, but not have (or have time to make) the appropriate inference from that evidence, it is difficult, if not impossible, to admit this criterion and still stop short of the Actual Beliefs View.

provide a usable decision-guide for agents such as Alex in Altered Option cases. It appears, then, that the best strategy for advocates of the Reasonable Beliefs View is to shift to the ‘Mini-Idealist’ approach. According to this approach, what a reasonable person in circumstances such as Daria’s (circumstances which include not having investigated at all) would believe depends on what she would have believed if she had taken advantage of her *last possible opportunity* to investigate.

In describing Daria’s case I said that her last opportunity to investigate is presented on November 2, when she could read the local newspaper’s editorial column on the election. But strictly speaking, this is not correct. On November 3, when she is standing in the voting booth, she has one final opportunity to investigate—but in this case, the kind of investigation in question is *deliberation*, rather than research. Although she has not investigated the candidates by listening to news sources, she does have some impressions about their characteristics, and she has the opportunity to review those impressions and evaluate, based on these impressions, which candidate seems best qualified to be governor. In other words, she could review the evidence available to her. Let us suppose such a review would lead her to believe that Aviles is the best candidate. Thus the Mini-Idealist approach should say that what it is reasonable for her to believe, as she decides for which candidate to vote, is what she would believe if she had taken advantage of this *very last* opportunity to investigate. Since, on the Mini-Idealist approach, it would be reasonable for her to believe that Aviles is the best candidate, it is subjectively obligatory for her to vote for him.²⁴

Thus, on the Mini-Idealist approach the Reasonable Beliefs View construes what an agent would be reasonable to believe (in the kinds of cases we are considering) to be the belief she would have arrived at had she deliberated as she ought to have done at the very last moment before the decision. Possible advance research drops out of the picture, so long as the agent failed to pursue it.²⁵

²⁴ Here again I am temporarily setting aside questions of whether Daria would have any belief about what she would believe if she had taken advantage of her very last opportunity to investigate.

²⁵ What should we say in a different version of *Election*, in which Daria goes above and beyond the call of duty by investigating *all* of these news sources? Suppose in such a case that she would conclude that a fourth candidate, Diamond, is best. I assume in the text that she is only obligated to investigate one information source. But in this version of *Election*, Daria

What does the Mini-Idealist approach imply for *Space Medicine*? It implies that what it would be reasonable for Alex to believe at the time he has to medicate the space-station crew is what he would believe if he had (just now) deliberated about which antibiotic would be best. This removes the problem of his subjectively obligatory action being one he believes not to be available, since after deliberating presumably his subjective obligation would be to choose among the antibiotics that he believes are available. So, shifting from the Maxi-Idealist approach to the Mini-Idealist approach seems to rescue the Reasonable Beliefs View from the threat posed by the fact in Altered Option cases that research can remove some of the options that would otherwise have been available to the agent.

But this is moving too fast. Although it would be less common, it is also possible for *deliberation* to alter the options that are available to the agent at the time of action, since deliberation can also affect an agent's options. Consider the following example, in which the form of deliberation in which the agent needs to engage is calculation.

Psych experiment. Amy is the subject of a psychology experiment which uses an advanced form of functional magnetic resonance imaging (fMRI) of the brain's operations. She has been outfitted with the fMRI headset, and is told by the experimenter that the computer screen in front of her will display arrays of options that describe rewards that she can receive. She must select one option in each array 20 seconds after the display begins. The options are described in such a way that calculating which option is best requires some mental effort. The experimenter tells Amy that if she deliberates about the options with which she is presented, she will receive whichever option she selects. However, if she does not deliberate, the options actually available to her will be randomly changed from those first displayed, and the display will change to show this alteration just before the 20 seconds have elapsed. In this event she must choose without further time to deliberate. The experimenter will know from the fMRI signals whether or not she is deliberating once

would arrive at a different belief if she went above and beyond the call of duty in her investigations. This opens an interesting question (which I shall not pursue here) about what we should say a reasonable agent would believe in such a situation—and whether the belief of the supererogatory Daria counts as 'reasonable'.

the array of options is displayed. Amy's best strategy is to deliberate before each choice. The first array of options A, B, C, D, and E is presented to Amy. If she deliberates, she will conclude that option B is best. However, she fails to deliberate, and now must instantly select an option from the altered display which includes only options A, C, and D.

What option would it be reasonable for Amy to believe is best, given that she has not deliberated? According to the Mini-Idealist approach, what it would be reasonable for her to believe is what she would have believed if she had taken her last opportunity to deliberate. Let us suppose that if she had done this she would have concluded that option B is best. However, since she has not deliberated, option B is no longer available to her. Let us add to the case the experimenter's now saying to Amy: 'If you had deliberated, you would have concluded option B to be best.' Then the Reasonable Beliefs View tells her that she subjectively ought to select option B. But, since she did not deliberate, option B is not available to her, and the Reasonable Beliefs View has once again—as in *Space Medicine*—delivered a recommendation for action that the agent knows she cannot follow.

Of course, *Psych experiment* is an unusual example. But since deliberation takes time, and involves mental activity and effort, engaging in deliberation can and does make a difference to what options are available to agents as compared with the options they would have if they failed to deliberate. Since deliberation takes time, options can disappear with the passage of time (think of an agent playing a fast-moving computer game such as Tetris). Deliberation can interfere with an agent's ability to carry out other kinds of mental processes (think of a pianist who, if he deliberates about how to phrase a musical passage, will be unable to play all its notes accurately; or think of a job candidate who, if she deliberates about what to say in her forthcoming interview, will have a car accident en route to the interview). These cases may not be common, but nonetheless they occur. What they show is that switching from the Maxi-Idealist approach to the Mini-Idealist approach for defining 'reasonable beliefs' does not, despite our initial surmise, rescue the Reasonable Beliefs approach from the charge in Altered Options cases that it can deliver prescriptions for actions that the agent knows she cannot carry out.

Switching to the Mini-Idealist approach fails to rescue the Reasonable Beliefs View from the problem of recommending actions that are not available to the agent. It also reveals another problem for the Reasonable Beliefs View. Consider Amy in *Psych experiment*. In each round of the experiment she has 20 seconds to deliberate about which option to select. Let us now add the detail that there are various *types* of deliberation in which she can usefully engage. If she uses the full 20 seconds, then a deliberation of type K (perhaps calculating the expected values of her options) would be best. If she has only 10 seconds available, then a deliberation of type L (perhaps calculating the maximum possible loss for each option) would be best. If she has only 5 seconds available, then a deliberation of type M would be best. If she has only 2 seconds available, then a deliberation of type N would be best. If she has only 1 second a deliberation of type O would be best, while if she has only $\frac{1}{2}$ a second then a deliberation of type P would be best. If she has less than $\frac{1}{2}$ a second, the best she can do is to use a deliberation of type Q—choosing the option that strikes her as most advantageous.

Suppose, as before, that Amy does not use any of K, L, M, N, O, or P, but now must choose an option with less than $\frac{1}{2}$ a second left. According to the Mini-Idealist approach, what would be subjectively obligatory for Amy to do is the action that would be best in light of the belief that it would be reasonable for her to have in the circumstances: that is, the belief she would have if she had utilized her last opportunity to deliberate. But what exactly is her *last opportunity* to deliberate? Surely it is not the opportunity which she had when 1 second was left, since she had another opportunity: namely, the opportunity to use method P when only $\frac{1}{2}$ a second was left. But why is *this* her last opportunity, when she still had the opportunity, when less than $\frac{1}{2}$ a second was left, to use method Q: that is, choosing the option that strikes her as most advantageous? But is method Q really a method of *deliberation*? At this point it seems completely unclear where one would draw the line between Amy's 'last opportunity to deliberate' and her simply selecting the option that strikes her as best. There seems to be no theoretical ground for distinguishing these. Thus it appears as though, if we take the Mini-Idealist approach seriously, we are driven away from the Reasonable Beliefs View into the arms of the Actual Beliefs View, since Amy's last opportunity is to use method Q (choosing the option that strikes her as most advantageous)—but method Q just *is* making

the best choice in light of her actual beliefs at the time of choice. Although *Psych experiment* is a somewhat artificially structured case, in fact many—perhaps most—real cases involve exactly the same phenomenon as the one highlighted here: even when the possibility of research has passed and we are reduced to mere deliberation prior to some decision, what counts as deliberation shrinks as we approach closer and closer to the decision itself, until finally there is no distinction between deliberation and choosing what is best in light of one's actual beliefs. To hold that the subjectively obligatory act is the one that is best relative to one's evidence does not take into account these facts about deliberation *about* that evidence.

IV. CONCLUSION

We have been investigating whether what is subjectively right for an agent depends on her actual beliefs or on what it would be reasonable for her to believe. I have argued that neither of two canvassed arguments offered in favor of the Reasonable Beliefs View holds up under scrutiny, and further argued, more decisively, that this View is subject to fatal problems. The point of principles of subjective moral status is to provide decision rules that can be used by any agent to make her decision, no matter how impoverished or distorted her beliefs about her options may be. Since agents may often have no beliefs, or be uncertain, about what they would believe if they had investigated as they ought to have done, it appears that the Reasonable Beliefs View cannot meet this demand. Reflection reveals that principles of subjective obligation based on the Reasonable Beliefs View are particularly bad at dealing with situations in which whether or how an agent investigates affects not just what beliefs she would have, but also what opportunities she would actually confront at the time of decision. In this type of case, if the Reasonable Beliefs View is spelled out in terms of the Maxi-Idealist approach, it will often prescribe actions that the agent knows are no longer available, and so offer advice which she cannot follow. On the other hand, if the View is spelled out in terms of the Mini-Idealist approach, the Reasonable Beliefs View remains vulnerable to this problem, and in addition collapses into the Actual Beliefs View for agents who fail to deliberate until the very last moment.

My conclusion is that we should resist the temptation to define what is subjectively right in terms of beliefs that the reasonable agent would have had, and instead define this notion in terms of the beliefs that the agent—reasonably or not—actually has at the time of decision. Just as you must go to war with the army you have rather than the army you might wish to have, you must make moral decisions with the beliefs you actually have, not the beliefs you might wish to have.²⁶

²⁶ I am grateful to Alvin Goldman, Doug Husak, Elaine Leventhal, and Evan Williams, as well as to the participants of the 2010 Arizona Workshop on Normative Ethics and the 2010 American Philosophical Association Pacific Division session on ‘Non-Ideal Theory’, for very helpful discussion and comments on earlier versions of this paper.