



OBJECTIVISM AND PROSPECTIVISM ABOUT RIGHTNESS

BY ELINOR MASON

JOURNAL OF ETHICS & SOCIAL PHILOSOPHY

VOL. 7, No. 2 | MARCH 2013

URL: WWW.JESP.ORG

COPYRIGHT © ELINOR MASON 2013

Objectivism and Prospectivism About Rightness

Elinor Mason

IMAGINE THAT I AM IN MY CAR, approaching a junction I know well. It is the middle of the night, there are rarely other cars and the chances that I would see a car coming if there was one are pretty good. Of course, there is a small chance that there is another car coming, and that if I speed through I will not see it, and, in that case, at the speed I am going, a crash would probably involve fatalities. On the other hand, if I get home more quickly, I can send an important email before the London stock exchange opens, and potentially save a lot of people a lot of money. Is it morally permissible for me to speed through the junction without stopping? What if we stipulate that there is no car coming? I cannot know it, but let us say that speeding through would in fact be perfectly safe. The question is about how rightness is determined. Is it determined by what is actually the case, or by what I believe to be the case?

A longstanding debate within moral theory, and consequentialism especially, concerns the extent to which rightness depends on the agent's epistemic and other limitations. Objectivists claim that the right action is the one that actually would have the best consequences.¹ Prospectivists claim that the right action is the one that is prospectively best – for now, let us gloss that as the one that it would be rational to choose given the agent's unavoidable epistemic limitations.² Subjectivism is usually formulated in terms of what the agent actually believes would be morally appropriate.³ My aim in this paper is to explore what is at issue between proponents of objectivism and proponents of prospectivism. I do not discuss subjectivism here; for reasons that will become apparent, we should think of subjective rightness as being in a rather different category to prospective and objective rightness. I argue that prospectivism is the correct account of moral rightness. Morality requires that, where possible, we build uncertainty into our ranking of options.

1. The Regan/Jackson Example

Recent debate in the literature is based around an example that appears to show that prospectivism must be correct. Donald Regan, Frank Jackson

¹ Defenders of objectivism include Moore (1912), Ross (1930), Feldman (1986, 2003), Driver (2002, 2012) and Graham (2010).

² Defenders of prospectivism include Smart (1973), Gruzalski (1981), Jackson (1991), Oddie and Menzies (1992), Hooker (2000), Timmons (2002), Howard-Snyder (1997) and Zimmerman (2006, 2008). (I have taken the term “prospectivism” from Zimmerman.) Various degrees of objectivity are possible – I explain below.

³ For example, Zimmerman (2008) (arguing against the account) and Holly Smith, who develops a detailed account of subjective obligation in a series of papers – see particularly (2010). Prichard (1932) and Ross (1939) (he changed his mind in favor of subjectivism after being convinced by Prichard) both argue for subjectivism, though not about consequentialism in particular.

and others have presented versions of this example. Jackson's version is as follows:

Jill is a physician who has to decide on the correct treatment for her patient, John, who has a minor but not trivial skin complaint. She has three drugs to choose from: drug A, drug B, and drug C. Careful consideration of the literature has led her to the following opinions. Drug A is very likely to relieve the condition but will not completely cure it. One of drugs B and C will completely cure the skin condition; the other though will kill the patient, and there is no way that she can tell which of the two is the perfect cure and which the killer drug. What should Jill do?⁴

Let us stipulate that, in fact, drug B will cure the condition. Then, according to the standard conception, it would be objectively right for Jill to prescribe B. Jackson's point is that this is irrelevant – clearly she ought to prescribe drug A. She should not be trying to do what is objectively right and she should not be trying to do what is most likely to be objectively right; rather, she should do what she knows is objectively wrong: She should prescribe the safe drug.

If we are convinced by Jackson's claim that Jill ought to prescribe the safe drug, we have a strong argument against objectivism. Objectivism tells us that it would be wrong to prescribe the safe drug, and so apparently must disagree that Jill should prescribe it.⁵ The prospective view says that Jill ought to do what is prospectively best, where that is usually what has highest expected utility.⁶ Going prospective solves the problem raised by Jackson's example – what we ought to do is prescribe the safe drug, and that is also what is right according to prospectivism.

It is not completely clear how the argument for prospectivism works here. On the one hand, it is simply an appeal to an example: Prospectivism gives the intuitively right answer. However, hardheaded objectivists remain unconvinced. Fred Feldman, for example, describes a case with the same structure as Jackson's, but insists that the right thing to do is

⁴ Jackson (1991: 462-463), see also Regan (1980: 264-5, n. 1). Fred Feldman discusses a version of this example in Feldman (1986: 46-47), but Feldman's conclusion remains that what you ought to do is prescribe the drug that would actually be best. More recently Parfit gives a version of the example (2011: 159). Parfit's overall conclusion is that there are many senses of right and wrong and that they are all useful.

⁵ Perhaps objectivism can tell the agent that although prescribing drug A is wrong it is nonetheless what she ought to do. I will come back to that possibility.

⁶ The term "prospectively best" leaves open the possibility that what it is most rational or moral to do with respect to uncertainty may not always be what has the highest expected utility. First, as is well rehearsed in the literature on consequentialism, there may be cases where calculating expected utility is self-defeating because it is too difficult or too time consuming, or the wrong sort of thought process to be having. Sometimes it is better to deal with uncertainty with some sort of heuristic device. Second, it may be that not even ideally rational agents would always deal with uncertainty by means of expected utility calculations. Perhaps sometimes maxi-min is a more morally appropriate strategy (for example, when risks fall disproportionately on one person or one group). Perhaps sometimes (for example where the values in play are those that seem vulnerable to being undermined by calculation) we should act according to preformed dispositions, or our gut feelings.

what is objectively best.⁷ The example on its own does not constitute an argument.

Alternatively, the argument might have something to do with “action guidance.” However, there are two problems with this. First, the notion of action guidance is not at all clear. Second, on most understandings of the notion of action guidance, prospectivism is not action guiding.

Finally, the argument might be that the example somehow shows that prospectivism rather than objectivism captures the central or primary use of “ought.” In what follows, I will argue that none of these are good argumentative strategies. I will then present my own argument for prospectivism.

Of course, our use of the various terms here is not univocal. In fact, one common strategy in response to the Jackson example is to say that we have various uses of the words “ought” and “right,” we use them in objective, prospective and subjective senses, and that there is not much more to say. To some extent I am sympathetic to this position. My own view is that we use the words “ought” and “right” and other corollaries for different notions depending on what we are interested in. Sometimes we are interested in what actually would have happened, because we want to learn from our mistakes and do better in the future. So it is natural for Jill to say, “Aha, blood type determines which drug is the cure. I see now the right thing would have been to prescribe drug B.” At other times we might be more interested in a retrospective agent evaluation. One sort of agent evaluation focuses on whether what an agent did was rational. Presumably, we are keen that our doctors behave rationally, and so it would be natural for the patient to say, “Although it turns out that B would have cured me, Jill did the right thing in prescribing the safe drug, and it would have been wrong for her to prescribe B.” And sometimes we are interested in a different sort of agent evaluation. We may be interested in the agent’s conscience, and so even when the agent has acted in ways that to the rest of us seem irrational, we might say, “He has a misguided view of what his duty is (he has been brainwashed by years in the army) but he is a man of conscience and when he ordered the court martial he really thought he was doing his duty. In that sense he did the right thing.”

It is true that common usage ranges over various senses of “ought” and “right.” Nevertheless, I think there is work to be done in clarifying the different notions, and once we have done that we can say more about exactly why each of the notions is interesting. And I shall argue that prospectivism is the best account of moral rightness.

2. Action Guidance

It is often suggested, for example by Jackson himself, that prospectivism is action guiding whereas objectivism is not. Jackson says: “[T]he fact that a course of action would have the best results is not in itself a guide to

⁷ Feldman (1986: 46-47). Recently, Feldman has developed an account of subjective rightness as a supplement to his objectivism (2012).

action, for a guide to action must in some appropriate sense be present to the agent's mind. We need, if you like, a story from the inside of an agent to be part of any theory which is properly a theory in ethics."⁸

Let us start with a brief exploration of the notion of "action guidance." There is an ambiguity in the notion of action guidance, and a parallel ambiguity in the notion of prospectivism, so it is worth clarifying what is going on. There seem to be various things that we might mean when we talk about action guidance. At one end of a sliding scale, we might be asking for an absolutely foolproof instruction: "Put one foot in front of the other until you reach the door, raise your arm and grasp the door handle, turn the handle," and so on. Of course, even that requires some knowledge and inference, some "practical wisdom" in the weakest sense of that phrase. At the other end of the scale are instructions that would be action guiding for very practically wise and rational agents, but not for most of us. "Do the thing that correctly balances magnanimity and justice," for example, is not very helpful for most agents. What are we asking for, then, when we ask that a theory be action guiding?⁹

On the one hand, we could take action guidance as a vague but objective notion. We do not call an instruction action guiding if only a fully rational, practically wise agent would understand it, but, on the other hand, we do not say that an instruction is not action guiding just because a very few, very limited agents do not understand it. Somewhere in the middle we say that an instruction is action guiding, in which case action guidance is an objective notion, in that it does not depend on the understanding of particular agents, but rather on the understanding of a somewhat idealized agent. This seems close to the way that we use the notion of action guidance, but as we will see below it renders the notion useless in the debate between various versions of prospectivism.

Alternatively, we might insist on a more subjective account of action guidance, according to which an instruction is action guiding for P if P can actually understand and follow the instruction. So the detail in an action guiding instruction will vary according to P's cognitive and psychological limitations. Some agents need more help than others. This subjective version of the notion of action guidance is the one that is being tacitly appealed to in defenses of subjectivism about rightness. The thought is that there is no point in telling an agent to do what a rational person would do – this is not an instruction that an imperfectly rational person can cash out. Underlying the subjective use of the notion of action guidance is a worry about responsibility: How can we be responsible for not following an instruction that we could not follow?

Just as we can formulate more and less objective forms of the notion of action guidance, so prospectivism can be more or less objective. There are various elements in prospective bestness, and we need to be clear

⁸ 1991: 466-467.

⁹ Thanks to Julia Annas for pressing me on this.

about whether we have an objective or subjective account of each.¹⁰ First, there are the beliefs that the agent has about her situation – should we take the beliefs the agent actually has, or should we take the beliefs that it would be rational for an agent in her position to have? Second, there are the probability assignments – should we take an agent’s actual probability assignments, or the probability assignments that it would be rational for her to make? Finally, there is value itself. Should we take the agent’s actual value system, the value system that it would be reasonable for her to have (which might vary over time and cultures) or the value system that is in fact correct (if there is such a thing)? Further, there can be more or less objective accounts of what it is reasonable or rational for an agent to believe, and a thorough taxonomy of possible views here should recognize that variation, too.

Possible accounts of what prospective bestness is are parallel to possible accounts of what action guidance is. An extremely objective prospectivism uses the beliefs, probability assignments and value system that it would be rational (or, if applicable, correct) for the agent to have. Fully subjective prospectivism uses the agent’s actual beliefs, probabilities and values, whatever they are. A moderately objective account of prospectivism involves a slightly vague position somewhere between the two extremes.

The most compelling account of objective prospectivism would not be extremely objective. Full rationality is too high a standard to capture what we are interested in here. The point of prospectivism is to give an account of rightness that takes uncertainty into consideration. Rightness, like its close relatives in epistemology, knowledge and justification, must

¹⁰ Neither Jackson nor Regan discusses this issue. Jackson does put his cards on the table (1991, 1986). Jackson’s account of the prospective ought is partly objective. First, Jackson uses an objective account of value. He describes the values that figure in the expected utility calculation in terms of idealized desires:

We can think of consequentialism’s value function as telling us what, according to consequentialism, we ought to desire. For a person’s desires can be represented – with, of course, a fair degree of idealization – by a preference function which ranks states of affairs in terms of how much the person would like the state of affairs to obtain, and we can think of consequentialism as saying that the desires a person ought to have are those which would be represented by a preference function which coincided with consequentialism’s value function (1991: 464).

See also Jackson (1986: 352). Second, although Jackson is explicitly committed to using actual beliefs as opposed to idealized beliefs (464), it is not clear what he thinks about probabilities. The phrase “subjective probability function” is ambiguous. It is usually used to mean the probability that one would rationally assign given one’s actual epistemic position (it is subjective because there are no constraints on how the agent reached that position). But it could be used to mean the probability that one actually assigns. I am not certain what Jackson intends. Oddie and Menzies (1992) distinguish between two forms of prospectivism (though their terminology differs from mine): “subjectivism,” which uses subjective probabilities, and objectivism which uses objective probabilities (514-515). They do not explicitly recognize that there can also be levels of subjectivism about value.

apply to real agents even if it is not fully relativized to real agents.¹¹ If a fully rational agent is the standard of rightness or justification, it is not clear that the standard is applicable to real agents – first, it is just intuitively too demanding, but second, problems arise because fully rational agents would never end up in the situations real agents find themselves in, and so it is not clear that it even makes sense to talk of what fully rational agents would do or believe in those situations.

Objective prospectivism should instead allow some vagueness, and hold that the standard of right action depends on what a reasonable agent would believe in the circumstances. A reasonable agent is not one with exceptional powers of rationality, but one who is rational enough, good enough at making probability estimates, good enough at knowing what the values at play in a situation are. Of course this is not to say that actual prospectivists accept this account. Jackson, for example, appears to be committed to a version of prospectivism according to which rightness depends on actual beliefs, reasonable probability estimates and the true account of value. Basing rightness on actual beliefs, as opposed to the beliefs a reasonable agent would have, has the disadvantage that these beliefs could be crazy. In other words, it has the disadvantages of subjectivism, which I will come back to. In this paper I will be discussing “moderate objective prospectivism”: reasonable beliefs, reasonable probability estimates and a reasonable understanding of value.¹²

Moderate objective prospectivism is not necessarily action guiding. What a reasonable agent would do given a certain epistemic position is not a guide to action for a less-than-reasonable agent – it is not a “story from the inside.” Of course, if we are assuming a moderately objective account of action guidance, we can say that an instruction is action guiding even though it is not action guiding for a particular agent. However, in that case the appeal to action guidance does not provide any independent support for an objective version of prospectivism. Someone who is primarily interested in a story from the inside, and the related issues of responsibility and action ownership, should simply reject an objective account of action guidance. If we take a more subjective view of action guidance, we must admit that much of the time, normal agents will not be able to derive a plan of action from instructions given to them by moderate objective prospectivism. This is what seems crucial – not the fact that, given our definition of action guidance, we can insist that the theory is “action guiding.”¹³

¹¹ For an interesting discussion of the relationship between objectivism and subjectivism in ethics and in epistemology, see Richard Feldman (1988).

¹² Oddie and Menzies (1992) offer a detailed and technical account of objective prospectivism. My account differs from theirs in that I explicitly include uncertainty about other facts and values; they focus exclusively on uncertainty about probability.

¹³ For a very vivid illustration of how prospective consequentialism can fail to be action guiding, see Fred Feldman (2006). Prospectivists such as Oddie and Menzies (1992: 517) and Zimmerman (2008: 71) acknowledge that the motivation for prospectivism is that the theory be action guiding.

Given, then, that standard versions of prospectivism are not action guiding, should we be concerned about this? Holly Smith raises the question vividly by showing that prospectivism is itself vulnerable to a Jackson-style counterexample.¹⁴ In Smith's example, the doctor, Harry, has a choice between three courses of action. He knows that all of them have risks, but also that all of them have the potential to cure. He does not know any of the probabilities, but he does know what the worst possible outcome is in each case: losing the use of a foot (E); losing the use of the left hand, which would be worse than that (F); and losing the use of the right hand, which would be worse still (G). A colleague tells him that option E is not prospectively best.

The example does not work if we are imagining a version of prospectivism according to which the prospective best depends on the beliefs an agent actually has. In that case, the testimony received from Harry's colleague gives Harry some knowledge about what is prospectively best for his colleague, but it does not change what is prospectively best for him. The colleague, unlike Harry, knows what the probabilities of the bad outcomes are. Given that Harry still lacks the crucial information – which of F and G is better – E remains prospectively best for him. The colleague's testimony does not provide enough information to change subjective prospective bestness. Thus Jackson's own version of prospectivism does not seem vulnerable to this version of the counterexample.

What if, on the other hand, we take prospective bestness to depend on the beliefs a reasonable agent would have? In that case, what is subjectively prospectively best for Harry may not be what is prospectively best, and, unlike in the previous version of the case, it is conceivable that the colleague can give Harry information about what is prospectively best for Harry. It might be clearer to stick more closely to Jackson's original case. Imagine that Jill has a choice between three treatments. She knows that E will most likely result in the loss of a foot, and she knows that one of F and G will result in a cure and the other in death, but she does not know which is which. She knows that if she were more reasonable she would be able to figure out which of her options is prospectively best – the fact that she would be able to figure it out if she were reasonable is what distinguishes this case from the original case. Because Jill knows that she would be able to figure out which of F and G would kill her patient, she knows that F or G is prospectively best. Now it is clear that the information that E is not prospectively best is useless to her. She does not know whether F or G is prospectively best, and as one of them would be catastrophic, there is a clear (subjective) sense in which she should not risk choosing F or G. Thus Smith is right: There is a Jackson-style counterexample to objective prospectivism.¹⁵

¹⁴ Smith (2011). Smith's larger point is that the possibility of a Jackson-style example does not show that an account of rightness is flawed, and so it is not an objection to objectivism. Her own view involves "dual oughts," both subjective and objective, and she allows that they can conflict.

¹⁵ Thanks to Wlodek Rabinowicz for helping me to clarify my thinking on this issue.

However, proponents of objective prospectivism, just like proponents of objectivism before them, are quite justified in taking this as irrelevant. They can insist, as Feldman does about objectivism and Zimmerman does about objective prospectivism, that the right thing to do is nonetheless what is right according to the objective standard. This may not be accessible to the agent – it is certainly not “a story from the inside” – but it is not irrelevant. To put it another way, both objectivism and objective prospectivism are concerned with what the reasons are, not with whether reasons are reasons for particular agents.

It might be objected that “ought” implies “can.” But there is certainly no clear violation of the “ought implies can” principle in either objectivism or objective prospectivism – in both cases the options are limited to those that an agent can physically do. If the “ought implies can” principle is being used in a richer sense than this, it cannot be taken as an argument in itself – it needs justification.¹⁶ The objectivist and the objective prospectivist are entitled to say that they are not interested in richer notions of “ought implies can” and the corresponding notions of responsibility, praise and blameworthiness, conscience and so on. This is the realm of subjective rightness. Objectivism and objective prospectivism specify goals, and the debate between them is about where the goal is – it is not relevant that some agents cannot figure out how to meet the goal.

Objectivism and moderate objective prospectivism are on all fours with respect to Jackson-style counterexamples. The fact that we think that there is something else an agent should do when they do not know what is objectively or prospectively best shows that we need a subjective “ought” – that there is one that we intuitively understand – an “ought” that depends on the agent’s actual position and takes into account her actual limitations. This “ought” needs to be elaborated. But I will leave that task for another time. My focus in this article is on objective prospectivism: Given that the argument for objective prospectivism over objectivism cannot be an argument based on action guidance, is there any argument to show that it is a superior account of an objective ought?

2.a. Action Guidance and Subjectivism

Before getting back to the main topic, I will digress briefly to point out that the subjective “ought” is not straightforward. It might seem obvious that the one virtue of subjectivism is that it is action guiding.¹⁷ However, consider a subjectivism such as that proposed by Ross. Ross proposes that the right thing for someone to set himself to do is what “he thinks to be morally most suitable in the circumstances as he takes them to be.”¹⁸

¹⁶ For defenses of subjective rightness based on an appeal to a rich notion of “ought implies can,” see Howard Snyder (1997) and Mason (2003).

¹⁷ Holly Smith (2011) and Fred Feldman (2012) both defend accounts of subjective rightness, where subjective rightness is intended as supplementary to objective rightness, and is needed because it is action guiding. These views are not fully subjective.

¹⁸ Ross (1939: 161).

Imagine someone sincerely asking for advice in a difficult moral situation. Such a person might say, “I do not know how things will turn out, though I have a good guess. But more worryingly, I am not sure whether keeping a promise is more important than producing good consequences – what should I do?” On Ross’ view, this person ought to do what they think is right. But they do not know what is right – that is the whole problem. Subjective rightness has no substance to it, and so, in an important sense, is not action guiding at all.

A fully subjective view is action guiding only in a very weak sense. It tells us to obey our conscience, but no more than that. In order for an instruction to be action guiding there must be some objective substance to it.

3. The Primary Notion

There is certainly ambiguity in our obligation terms. But there remains a question about which is the central or primary term. As Zimmerman puts it, there must be some “overall” moral obligation: “It is with overall moral obligation that the morally conscientious person is primarily concerned.”¹⁹ But what is the morally conscientious person concerned with? It depends on what we mean by the question. If we think about what the morally conscientious person will actually do, it seems that we cannot say more than that she will do what is subjectively right: She will try as hard as possible to make a rational assessment of the situation and act according to her own assessment of what is right. That is the most anyone can do. Why think more should be included in the description than that? Why think, in particular, that a morally conscientious person must actually get it right? Zimmerman’s own argument is that full subjectivism is absurd – it has the consequences that agents always know what they are supposed to do – they can simply introspect to find out; that an agent with no beliefs about what is best has no obligations; that morally misguided people, even severely misguided ones, like Hitler, are acting rightly; and finally that the subjective view may violate the “ought implies can” principle – one might believe that one can do things that are in fact impossible.²⁰

Subjectivism does seem absurd in some ways, but this does not, in itself, provide a good argument for defining the morally conscientious person as someone who is reasonable in their beliefs, probability estimates and account of value. All we can say about the morally conscientious person is that they will do what they think best. If we add that they are correct in their assessment of the situation, probabilities and values, we are just begging the question against subjectivism. So thinking about what the morally conscientious person would do does not help with the debate about the “primacy” of the various notions.

¹⁹ 2008: 2. Ewing (1947: 122-144) accepts that there are three valid senses of ought (corresponding to what I call the objective, the subjective and the prospective) but argues that prospectivism is the primary sense of “ought.”

²⁰ 2008: 11-12.

Clearly, this is not what Zimmerman intends. He is asking not, “What would the morally conscientious person do?” but something like, “What would the morally conscientious person ideally hope to do?” This question seems closer to a “central” or “primary” concept of rightness in that it anchors rightness to the standards of morality. We have a fairly deep intuition (as expressed by Zimmerman’s *reductio ad absurdum* of subjectivism) that rightness is idealized to some extent – doing the right thing must involve meeting an objective standard. But if the question we are asking is, “What would the morally conscientious person ideally hope to do?” then we have nothing that distinguishes prospectivism from objectivism. In one sense the morally conscientious person hopes to do what is actually best; in another sense she hopes to do what is rational. So the notion of the morally conscientious person is really no help at all here.

One recent argument for the primacy of objectivism, proposed by Peter Graham, is that when we get more information about a choice situation, we take it that we are getting information about what our obligations were all along. We do not think that our obligations have changed as a result of the new information.²¹ This is not a strong argument. The way that we speak often implies this, as Graham points out. But it is not hard to think of cases where the way that we speak favors the prospective or subjective account of rightness. Imagine that Jill finds out after the event that drug B would have cured her patient. It is perfectly natural for her to say that she should have prescribed drug B. It is equally natural for her to say that she did as she ought in prescribing drug B at the time. Is there any deep oddity in saying that our obligations change as our information changes? Graham argues as follows:

Promise: I have promised Kenneth that I will turn on a certain light at 6 p.m. All my evidence indicates that it is not 6 p.m. I ask you what I am morally obliged to do. You know of my promise and that it is 6 p.m., but you also know that I neither believe, nor have any evidence, that it is 6 p.m. You tell me that I am morally obliged to turn on the light.

Here is an argument:

1. In Promise, you speak truly in saying I am morally obliged to turn on the light.
2. According to no plausible subjectivist moral theory do I have a moral obligation to turn on the light in Promise.
3. Therefore, objectivism is true.

Premise 1 is intuitive and premise 2 follows from the description of the case.²²

²¹ Graham (2010: 91-92). As Graham points out, Zimmerman’s answer to the puzzle – that when someone else tells you that you ought to do Φ based on information that they have and you do not, they are lying, albeit morally obligatorily lying – is not very satisfactory. But we do not need to say that someone is lying – they are just summarizing their own epistemic position and if they are reliable that is evidence that you should question yours.

²² Op. cit., 92.

Graham says that premise 1 is intuitive. In fact, premise 1 just begs the question. According to subjectivism (and some semi-objective prospectivisms) someone who says that you ought to do Φ , when that is not indicated by your own evidence, is obviously not saying something that is straightforwardly true. Of course, it would be a natural thing for someone to say, and the subjectivist needs an account of what is going on in cases like Graham's Promise.²³ There are various possibilities. Zimmerman's suggestion, that the advisor is lying in a morally obligatory way, strikes me as clumsy, but something close to that must be right. Here is a better suggestion. Someone who tells you that you ought to do Φ (when your evidence, or the evidence as you see it, does not indicate that) is saying something that is elliptical for, "The evidence from my point of view points to doing Φ ." And that is what we do in fact understand in these situations. Imagine that Jill's reliable and knowledgeable colleague comes into her office as she is puzzling over what to do. "You should prescribe B!" says the colleague. What would Jill say? She would obviously not say, "No, that is false. Given my evidence I should prescribe A." On the other hand, she would not say, "Whatever you say, wise one." She would of course react to the statement as a statement about the evidence, and would presumably ask her colleague what the evidence was, thus coming to the same evidential position as her colleague.

Julia Driver offers another argument for the primacy of the objective "ought." Driver argues that even though you ought to do what is prospectively best in cases like the Jackson example, that is somehow dependent on the objective best.²⁴ But what could "dependent" mean here? The Jackson example shows that what is prospectively best is not straightforwardly related to what is best – doing what is prospectively best is not a way of doing what is best, or getting as close as possible, or approximating, or in any other way aiming for what is best. Driver suggests that we need an explanation of why it would be better if we did have more information in these cases. But we do not need more of an explanation than we already have – the outcome would be better just because it has more goodness. All forms of "ought" in such examples depend on goodness. The objective "ought" picks out which of the agent's options would have the most goodness. The objective prospective "ought" picks out the option that is the most rational compromise of goodness and risk.

I turn now to my own argument for prospectivism, which rests on the distinction between rightness and goodness.

²³ For more discussion of this issue, see Kolodny and MacFarlane (2010), and Keiseweter (2011).

²⁴ Driver (2012; 2013: 19-22).

4. Rightness and Goodness

The point of Jackson's example is that it is obvious that Jill morally ought to choose the option that is prospectively best. But why is this so obvious? Intuitively, it is because she should not take the substantial risk of a really terrible outcome, the patient's death. In other words, risk has found its way into rightness. This is the crucial point in the Regan/Jackson example. Morality is not silent about how to deal with uncertainty. On the contrary, as the wide agreement about what Jill should do in Jackson's case shows, morality, on occasion, gives us very clear guidance.²⁵ It is clear what Jill should do, as almost all commentators on Jackson-style examples agree. It turns out that morality requires that we build uncertainty into our ranking of options. This is a complex point, especially given the divergence between different accounts of rightness and goodness and the different structures of deontological and consequentialist theories. My aim is to elucidate and defend the claim that rightness builds in risk.

There are two ways that moral theories can rank options: in terms of rightness (ought-to-be-doneness) and in terms of goodness. The terms "rightness" and "goodness" here are not perfect. For a start, rightness suggests a binary division between right and wrong, but I do not wish to inherit that view. "Ought-to-be-doneness" is perhaps more felicitous, as it suggests the possibility of a sliding scale – some things ought to be done more than others.²⁶ A ranking could include supererogatory acts and it might include suberogatory acts, or a ranking may simply be a hierarchy from best to worst with no markers of rightness or wrongness, or super- or suberogation.²⁷ In what follows I use "rightness" and "ought-to-be-doneness" interchangeably.

Goodness, unlike rightness, is independent of the agents who have brought about or might bring about the situation. Goodness resides in states of affairs. Not all moral theories have an account of goodness. A pure deontological theory could get by with rankings in terms of rightness only (though this is rare). By contrast, consequentialist theories prioritize the good. For a consequentialist, ought-to-be-doneness is entirely dependent on rankings in terms of goodness. Most theories, and perhaps the most plausible theories, are mixed: There are some principles of rightness that are independent of goodness, and some that depend on goodness. I am leaving it entirely open what goodness is, and of course that is not an easy question – traditionally it is something like happiness

²⁵ Holly Smith seems to think that that no moral theory ever has anything to say directly about risk (2011: 68). She thinks that whenever we talk about what to do in a situation of uncertainty we have switched to talking about a subjective principle. This is exactly the view that I am arguing against – objective rightness sometimes includes risk assessments. Smith dismisses the view that I am calling prospectivism as an unsatisfactory version of subjectivism – unsatisfactory because it is not action guiding (2011: 79-80). I discuss Smith's positive view in more detail elsewhere.

²⁶ See Norcross (2006).

²⁷ See Driver (1992) on suberogation.

or pleasure, but it may be something more complex, and it may even include deontological elements.²⁸

4.a. Deontological rightness and uncertainty

I will begin by considering some examples of pure deontological rightness – principles where there is no appeal to goodness at all. My aim here is not to defend, nor to attack the absolutist deontological position I describe. In fact I think it is a view held by very few. Perhaps in places Kant himself is an absolutist, but most contemporary deontologists are not absolutist: They allow that considerations of the good can be relevant, too. I am simply describing an absolutist deontological position and arguing that it (by contrast with principles based on the good) does not and should not be expected to build uncertainty into rightness.

Many familiar deontological principles avoid uncertainty by appealing only to what is clearly under an agent's control, for example, the instruction, "Do not lie." A lie is by definition intentional, so an agent must know whether she is lying so she can obey or not obey the principle. There is no luck involved; it is entirely up to her.²⁹ However, some deontological principles are less clear. Take, "Honor thy father and mother." Success in honoring could be entirely down to the agent's sincere effort at honoring, or it could be something that the agent could be in a position of uncertainty about. An agent may not know which attitude constitutes honoring, and so an analogue of the Regan/Jackson example could arise. An agent might know that attitude A would slightly dishonor her parents, whereas either B or C would honor her parents, and either B or C would dishonor them terribly. What does the deontologist say about this? Can the deontologist simply say that the agent should do what is prospectively best – namely, slightly dishonor her parents?

Absolutist deontology cannot appeal to a second-best when the agent is ignorant, because there is nothing that would make the allegedly second best option choiceworthy. It is not the right option (by stipulation) and it is not a good option, because according to absolutist deontology, goodness is not relevant – there is only rightness. So the theory leaves itself nothing to appeal to in cases in which the agent is not sure which option is right. The instruction to honor your parents exhausts the resources of absolutist deontology. The point is to honor your parents – if they are not honored, there is no second best, no compromise.

A deontological theory cannot be criticized for reaching the end of action guidance here. The demand for deontic principles to be action guiding goes only as far as the substance of the theory allows. Compare a set of instructions about making a Victoria Sponge cake. The instructions tell you to add white sugar, and might note that if you do not have white

²⁸ For a mixed account of the good, see Feldman (2004). See also the debate about "consequentializing," e.g. Brown (2011).

²⁹ It might be argued that we can be lying without knowing it because of unconscious forces, but this seems a little tenuous.

sugar you can use brown sugar. But the instructions will not tell you what to do if you do not have eggs, because you cannot make a Victoria Sponge cake without eggs. We do not demand action guidance when the aim of the action guidance becomes unattainable.

Deontology is in a slightly more peculiar position when it comes to principles that appear to depend partly on the good. Take, “Do not kill.” Surely at least part of the justification for such a principle is that death is bad for people – that states of affairs in which people are dead are worse than states of affairs in which they are alive. Clearly though, this is not the whole justification. Part of the point of “Do not kill” *qua* deontological principle is *what the agent does*. It does not matter whether the agent’s killing one person would save 10 – she should not kill. So, what should the deontologist say about a situation where the agent is uncertain about which action constitutes killing? Imagine that she is faced with three buttons: She knows that button A will kill one person, and one of buttons B and C will kill 1,000 and the other will kill none. It seems likely, again, that for an absolutist, the deontological instruction here must simply be, “Kill no one,” even though the agent cannot intentionally follow that instruction.³⁰ The agent cannot know what the right thing to do is – there is no more to be said.

This example illustrates a version of the classic disagreement between deontology and consequentialism. The deontologist keeps his hands clean; the consequentialist gets his hands dirty. But it illustrates something else, too – consequentialism can be action guiding when deontology cannot. Consequentialism can be action guiding because, for consequentialism, ought-to-be-doneness is entirely dependent on the good, and (as I shall argue) risk and uncertainty about outcomes can be built into rightness. But for a deontologist, the focus on the actor as opposed to the outcome means that uncertainty about outcomes is often irrelevant. This is why deontology seems so irrational to consequentialists when deontological instructions seem to be based partly on the good – why not take the good into account whenever it is at stake? Deontological principles like “Do not kill” seem to be an irrational *mélange* of concern about what the agent does and concern about what happens.

4.b. Goodness, Rightness and Uncertainty

Let us return, then, to deontic principles that are based partly or wholly on the good.³¹ Let us assume that in the Regan/Jackson case we are thinking in terms of principles based on the good. It is clear what would be best on any reasonable account of goodness – a cured patient. A dead

³⁰ See Jackson and Smith (2006) for an example of consequentialist bafflement at the inability of deontological theories to guide action. Jackson and Smith argue that all the possible ways that an absolutist might deal with uncertainty are flawed. My argument here is that the absolutist, by his own lights, does not have to deal with uncertainty.

³¹ Deontology can of course include some deontic principles that are wholly based on the good.

patient would be very bad, and an improved patient good, but not as good as a cured patient. So it is easy to see why it is tempting to say that what the agent ought to do is cure the patient. That is, after all, what would be best, and if ought-to-be-doneness is straightforwardly related to goodness, that seems the straightforward answer.

However, we should not assume that goodness and rightness are directly related. What is right can be relative to the position of the agent. “Honor thy mother and father” or “Save the drowning child” mean different things for different agents. An agent’s position can affect what is right for her. Thus it is at least a possibility that an agent’s epistemic position can affect what is right. When the uncertainty concerns which option is right according to an absolutist deontology, uncertainty does not affect rightness. So, for an absolute deontologist, instructions about what to do in the face of uncertainty can only be pragmatic. However, if rightness is based on goodness, and the uncertainty is about the amount of goodness that will be produced, we can, and indeed must (I shall argue), take that uncertainty into account in defining the right option. In other words, we should be prospectivists.

Unfortunately, there are various ways in which a theory could base rightness on goodness, and this complicates matters considerably. Let us first distinguish between two supposedly consequentialist views that base rightness on goodness. One, let’s call it “consequentialism,” bases rightness entirely on goodness. According to this view, the amount of goodness produced by your action is all that matters to rightness. The other, let’s call it “bestness fetishism,” is subtly different. It says that the largest possible amount of goodness is all that is relevant to rightness. To see the difference between these views, and in particular the difference between what they say about uncertainty, consider the following example:

Imagine that a tennis player has a choice between three strategies: She knows that if she plays A she will get a lot of points, but not enough to win; she knows that of the other two strategies, one will give her a resounding victory, but the other will result in a trouncing. Unfortunately, she does not know which of strategies B or C is the winning strategy.

A sport like tennis is analogous to bestness fetishism, because winning is all that matters. The player ought to go for strategy B or C – she ought to risk a trouncing in order to have a chance at winning. Getting a lot of points is not worth anything. Analogously, according to bestness fetishism, doing something reasonably good (like effecting a partial cure for Jill’s patient) is not worth anything – if curing the patient is possible, curing is the only worthwhile outcome. Clearly, consequentialism is not bestness fetishism (though some people confuse the two).³² Consequentialism is not bestness fetishism because for consequentialism, goodness is what matters, not bestness. And of course this is much more plausible as a moral theory than bestness fetishism – as our reaction to the Jackson

³² The basic consequentialist principle is often characterized as “Do what is best” or “Maximize the good.” These formulations are ambiguous.

case shows. Jill morally ought to give drug A, because the patient would be much better off partially cured than dead, and what really matters morally is how well off the patient is.

If goodness is what matters, and amounts of goodness can be ranked more or less continuously, then although more goodness is better than less, and of course maximal goodness would be best, it would be irrational to risk an outcome with minimal goodness in order to have a chance at maximal goodness.³³ An analogy might help here. Imagine a company that uses computer programs to make investments. The company's aim is to increase profits as much as possible, and betterness relates exactly to the amount of money made.³⁴ There is often uncertainty. The program must make "guesses" about which investments will do well and which will do badly. The software designers will design the programs so as to pick the prospectively best option. Of course, what counts as the prospectively best option will have to depend on data that has been entered by human agents. But the data entered is not what I am arguing about here; we can take it as given, just as we take Jackson's stipulation of the values and probabilities as given. The point is that if all that matters to the company directors is money, they will instruct their software designers to design programs that choose the prospectively best options. It would be irrational to design a program that would only act when it was certain of achieving the maximum gain possible in a situation, or to act so as to have the best chance of achieving the best possible outcome in any situation.

In sum, if goodness is what matters, then of course we should be prospectivists. Rightness is an instruction to an agent, and as such, is tailored to the agent as far as the aim of the instruction licenses. There is no inviolable rule that the instruction "Do Φ ," where doing Φ would be right must always be comprehensible and followable for an agent. Perhaps the agent falls below the reasonable level of rationality required to understand an instruction. Or perhaps the instruction is such that, when the agent lacks relevant knowledge, the theory has just reached the end of the road. This happens with absolutist deontology when the agent does not know how to do what is right, or does not know which option is right. But it does not happen when principles of rightness are based on goodness and the agent is uncertain about goodness. In those cases, the aim of the instruction "Promote the good" clearly licenses doing what is prospectively best: balancing potential good against risk of bad. And so the right thing to do in those cases is not something the agent need be

³³ There may be some discontinuities in moral goodness – for simplicity I am ignoring that possibility here. Where there are discontinuities the theory may not be able to build in uncertainty.

³⁴ In other words, there are no discontinuities in value – we are not taking shame or glory or other complex psychological effects into account.

uncertain about – she can figure out what is right by figuring out what is prospectively best.³⁵

Consequentialism is the simplest and most plausible theory that bases rightness on goodness. But most theories base rightness on goodness to some extent. In other words, even most deontological theories are moderate rather than absolutist. Whenever a principle of rightness is based on goodness, uncertainty about goodness can arise, and when it does, the theory has the resources to take that uncertainty into account in defining the right option.

4.c. Compromising with Rightness

My strategy has been to argue that when rightness is based on goodness, we ought to make compromises in order to minimize our losses as well as maximize our gains. But could we not make the same argument directly about rightness and wrongness? In a recent article, Peter Graham suggests that sometimes we ought to do something we know is wrong so as to avoid the risk of doing something that is even more wrong.³⁶ Prima facie this sounds odd – the phrase “we ought to do wrong” is infelicitous. We have these various notions – bestness, rightness and ought-to-be-doneness, as well as subjective notions. It would seem natural to keep “rightness” and “ought-to-be-doneness” together – the most natural idea is surely that if something is right you ought to do it. However, it may be that our ideas have outstripped our terms, and that it would be better to separate ought-to-be-doneness and rightness. If this is the case, we had better be very clear about what our revisionist meanings are.

Graham does not provide us with good reason to accept the infelicitous terminology. The first problem with Graham’s suggestion is that the theory itself does not provide the resources for a compromise in any clear way. Consequentialism relies on a more or less continuous ranking of outcomes in terms of goodness – what is prospectively best is just the rational (or reasonable) balance of good and bad outcomes and the chances they will come about. But of course, any theory can face risk: Bestness fetishists can be uncertain about which outcome will be best,

³⁵ Oddie and Menzies suggest the same argument (1992: 530), though they do not put it in terms of the distinction between rightness and goodness: “Those options with greater objective value are closer to the realization of the ideal than those options with less. As a result, a miss is not always as bad as a mile. Failing to achieve the ideal admits of degrees, and it is objectively preferable for the agent to perform an act with more objective value rather than less in those choices which do so fail.”

³⁶ Peter Graham (2010). This strategy is also suggested by Portmore (2011), Bykvist (2011), Driver (2012a). Smart suggests a similar approach: He says that what is right is what would actually produce best consequences (1973: 47) and yet that one ought to do what would maximize probable benefit (1973: 12), but Smart does not defend the divergence in concepts. Graham is defending objectivism over a view he calls “ability constrained evidence subjectivism” and refers to as “subjectivism” for short. In my terminology, the view that Graham is attacking is a moderate objective prospectivism – one according to which rightness depends on the evidence that is available (in an objective sense) to the agent.

and absolutist deontologists can be uncertain about which outcome will be right. The problem for bestness fetishism and for absolutist deontology is that it is not clear how a ranking should proceed after the top-ranked outcome. Take bestness fetishism: The right action is the one with maximal goodness. But let us say we do not know which that one is. For hardcore bestness fetishism, according to which the only thing that matters is the best outcome and all the others are worthless, the sensible thing to do is obviously whatever has the highest chance of being the best option. Likewise for absolutist deontology – if there is one right option and all the rest are (equally) wrong, then the sensible thing to do is whatever has the highest chance of being right.

But Graham is not talking about hardcore or absolutist versions of these views. He claims that it is not the case that all options other than the best are on a par – some are “less wrong” than others. So the question is, how do we rank options in terms of less wrongness? Consider the sort of case Graham has in mind. An agent has a choice between one option, A, which involves telling a lie, and two more, B and C, one of which will result in someone’s death, and one of which will have only good results. She does not know which of B and C will cause a death. Surely the moderate deontologist will say that although lying is bad it is not as bad as causing death, and so what she should do here is lie. In other words, surely rankings that could be put in terms of less wrongness could equally be put in terms of less badness.

In that case, we have the same answer that prospectivism gives (you ought to do what is prospectively best given the values at stake), but a different terminology. Whereas the prospectivist says that the prospectively best action ought to be done and is right, Graham says that the prospectively best action ought to be done even though it is wrong. So the question is, does anything really hang on saying that the prospectively best action is wrong, rather than what we all admit is not the best possible action in the circumstances? If rightness is to be separated from ought-to-be-doneness, we need a good reason – we need something that rightness is indicating is important and is not covered by the other terms we have. But we have the term “bestness” and, for both consequentialists and moderate deontologists, it seems that it is possible to say everything that needs to be said about the option that would in fact be best just by saying that it would be best.

What else could Graham mean? Let us try to imagine a ranking of wrong actions from really wrong to not so very wrong that does not boil down to a ranking of badness. Imagine that a deity has handed down a ranking of actions. There are two possibilities for how the deity has explained things: Either there is some sort of quantitative value attached to each action – whatever we call it, this is functionally equivalent to badness, compromise is possible (because we can see how to weigh the various options against each other and against uncertainty) and we should be prospectivists. Alternatively, there is a ranking with no quantitative values attached, in which case it is impossible to judge how to weigh the degrees

of wrongness against one another. Without knowing how much more wrong an alternative action is, we cannot judge whether we should risk it or not. So Graham's suggestion is either the same as mine (with a clumsier terminology) or theoretically unworkable.

5. Conclusion

There are many senses of “right” and “wrong,” “ought” and “ought-not.” We do use these words in different ways, but this is not to say that we cannot improve on our current understanding of them. My suggestion is not that we eliminate talk of objective rightness; my suggestion is that prospective rightness is superior as a conception of rightness. Talk of objective rightness could easily be replaced with talk of bestness. So understood, objective rightness, or bestness, is a perfectly useful notion – primarily useful for moral learning.

The argument of this paper has been that what we take for granted when talking about rationality – that the rational option varies with uncertainty – also applies to rightness. The right thing to do varies with uncertainty. Bestness is not all that there is to rightness.³⁷

Elinor Mason
University of Edinburgh
Department of Philosophy
elinor.mason@ed.ac.uk

³⁷ I would like to thank my colleagues in the Department of Philosophy at the University of Edinburgh for engaging with several versions of these ideas. I am also extremely grateful for invitations to speak, and for feedback, at: The Hebrew University in Jerusalem (Conference on Uncertainty organized by David Enoch, June 2011), and especially to Michael Zimmerman and Holly Smith for subsequent discussion; Symposium on the Work of Julia Driver at the University of Texas, San Antonio (Brackenridge Philosophy Symposium organized by Jill Hernandez, February 2012); Lund University Department of Philosophy, University of Arizona Department of Philosophy and the Society for Applied Philosophy, Oxford 2012. For reading earlier drafts and commenting, I would like to thank Stewart Cohen, Fred Feldman and Mark Timmons. Finally, thanks are due to two anonymous referees from this journal for valuable comments.

References

- Brown, C. (2011) "Consequentialize This," *Ethics* 121(4): 749-71.
- Bykvist, K. (2011) "How to Do Wrong Knowingly and Get Away with It," in *Neither/Nor. Philosophical Papers Dedicated to Erik Carlson on the Occasion of His Fiftieth Birthday*, Uppsala, Sweden: Department of Philosophy, Uppsala University.
- Driver, J. (2013) "What the Objective Standard Is Good For," in M. Timmons, ed., *Oxford Studies in Normative Ethics, Vol. 2*, New York: Oxford University Press.
- ____ (2011) *Consequentialism*, London: Routledge.
- ____ (2001) *Uneasy Virtue*, New York: Cambridge University Press.
- ____ (1992) "The Suberogatory," *Australasian Journal of Philosophy* 70(3): 286-95.
- Ewing, A. C. (1947) *The Definition of Good*, London: Routledge and Kegan Paul Ltd.
- Feldman, F. (2012) "True and Useful: On the Structure of a Two Level Normative Theory," *Utilitas* 24(2): 151-71.
- ____ (2006) "Actual Utility, the Objection from Impracticality, and the Move to Expected Utility," *Philosophical Studies* 129: 49-79.
- ____ (2004) *Pleasure and the Good Life: Concerning the Nature, Varieties and Plausibility of Hedonism*, Oxford: Clarendon Press.
- ____ (1986) *Doing the Best We Can*, Dordrecht, Holland: D. Reidel Publishing Company.
- Feldman, R. (1988) "Subjective and Objective Justification in Ethics and Epistemology," *The Monist* 71: 405-19.
- Graham, P. A. (2010) "In Defense of Objectivism About Moral Obligation," *Ethics* 121: 88-115.
- Gruzalski, B. (1981) "Foreseeable Consequence Utilitarianism," *Australasian Journal of Philosophy* 59: 163-76.
- Hooker, B. (2000) *Ideal Code, Real World*, Oxford: Oxford University Press.
- Howard-Snyder, F. (2005) "It's the Thought that Counts," *Utilitas* 17(3): 265-81.
- ____ (1997) "The Rejection of Objective Consequentialism," *Utilitas* 9 (2): 241-48.
- Jackson, F. (2001) "How Decision Theory Illuminates Assignments of Moral Responsibility," in N. Naffine, R. J. Owens and J. Williams, eds., *Intention in Law and Philosophy*, Aldershot: Ashgate, 19-36.
- ____ (1991) "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," *Ethics* 101(3): 461-82.
- ____ (1986) "A Probabilistic Approach to Moral Responsibility," *Studies in Logic and the Foundations of Mathematics* 114: 351-65.
- ____ and Smith, M. (2006) "Absolutist Moral Theories and Uncertainty," *The Journal of Philosophy* 103(6): 267.
- Kiesewetter, B. (2011) "'Ought' and the Perspective of the Agent," *Journal of Ethics and Social Philosophy*.
- Kolodny, N. and MacFarlane, J. (2010) "Ifs and Oughts," *The Journal of Philosophy* 107(3): 115-43.
- Lockhart, T. (2000) *Moral Uncertainty and Its Consequences*, Oxford: Oxford University Press.
- Mason, E. (2003) "Consequentialism and the 'Ought Implies Can' Principle," *American Philosophical Quarterly* 40: 319-31.
- Mele, A. (2003) "Agent's Abilities," *Nous* 37: 447-70.
- Moore, G. E. (1912) *Ethics*, London: Williams and Norgate.
- Norcross, A. (2006) "The Scalar Approach to Utilitarianism," in H. West, ed., *The Blackwell Guide to Mill's Utilitarianism*, Oxford: Oxford University Press, pp. 217-32.
- Oddie, G. and Menzies, P. (1992) "An Objectivist's Guide to Subjective Value," *Ethics* 102: 512-33.
- Parfit, D. (2011) *On What Matters, Vol. 1*, Oxford: Oxford University Press.
- Portmore, D. (2011) *Commonsense Consequentialism*, Oxford: Oxford University Press.
- Prichard, H. A. (2002/1932) "Duty and Ignorance of Fact," in *Moral Writings*, Oxford: Clarendon Press, pp. 84-101.
- Regan, D. (1980) *Utilitarianism and Co-operation*, Oxford: Clarendon Press.
- Ross, W. D. (1968/1939) *Foundations of Ethics*, Oxford: Clarendon Press.

- ____ (1930) *The Right and The Good*, Oxford: Oxford University Press.
- Sepielli, A. (2009) "What to Do When You Don't Know What to Do," *Oxford Studies in Metaethics* 4: 5-28.
- Smart, J. J. C. (1973) "An Outline of a System of Utilitarian Ethics," in J. J. C. Smart and B. Williams, *Utilitarianism: For and Against*, Cambridge: Cambridge University Press, pp. 3-74.
- Smith, H. M. (2011) "The 'Prospective View' of Obligation," *Journal of Ethics and Social Philosophy*.
- ____ (2010) "Subjective Rightness," *Social Philosophy and Policy* 27(2): 64-110.
- ____ (1991) "Deciding How to Decide: Is There a Regress Problem?" in M. Bacharach and S. Hurley, eds., *Essays in the Foundations of Decision Theory*, Oxford: Basil Blackwell, Inc., 194-219.
- ____ (1991) "Varieties of Moral Worth and Moral Credit," *Ethics* 101: 279-303.
- ____ (1989) "Two-Tier Moral Codes," *Social Philosophy & Policy* 7: 112-32.
- ____ (1988) "Making Moral Decisions," *Nous* 22: 89-108.
- ____ (1983) "Culpable Ignorance," *The Philosophical Review* 92: 543-71.
- Timmons, M. (2002) *Moral Theory: An Introduction*, New York and Oxford: Rowman & Littlefield.
- Zimmerman, M. J. (2008) *Living with Uncertainty*, Cambridge: Cambridge University Press.
- ____ (2006) "Is Moral Obligation Objective or Subjective?" *Utilitas* 18(4): 329-61.
- ____ (1997) "Moral Responsibility and Ignorance," *Ethics* 107(3): 410-26.