

Critical Notice

*Living with Uncertainty*

Clayton Littlejohn

cmlittlejohn@yahoo.com

In his previous work, *The Concept of Moral Obligation*, Zimmerman provided a careful and detailed analysis of overall moral obligation, conditional obligation, and *prima facie* obligation, dealt with a number of paradoxes of deontic logic, and argued we ought to be possibilists rather than actualists. (Possibilists and actualists disagree about the significance of future decision on present obligation, I shall have more to say about these views below.) He chose to bracket issues having to do with the agent's epistemic situation but said he was inclined to side with writers like G.E. Moore and Judith Thomson who thought that facts an agent is non-culpably ignorant of can have some bearing on the permissibility of an action.<sup>1</sup> This is a view he is no longer inclined to endorse. In *Living with Uncertainty*, Zimmerman defends the view that an agent's obligation is to do what is prospectively best whether or not that happens to be or believed to be objectively best.<sup>2</sup> After writing his earlier book, he was convinced by some of Don Regan's examples that an agent's obligations depend upon that agent's normative and non-normative evidence.<sup>3</sup> His latest work, like his earlier work, is an impressive achievement. It is packed with really interesting arguments and careful discussion of very complicated issues that results in a subtle and sophisticated account of moral obligation. It should be read by anyone with a serious interest in the concepts of moral obligation and moral responsibility. It is, like his previous work, a demanding read, but working through the details of the text is worth the effort.

To a rough first approximation, the prospectively best option is the option that maximizes expectable value given the agent's evidence. (Whereas *expected* value has to do with the actual values of outcomes and the probability that these outcomes will be brought about, *expectable* value takes account of both the probability that something will eventually occur and the probable value of such an occurrence where that value can vary between agents who assign all the same probabilities to all non-normative propositions but differ in the probabilities they assign to propositions about value given their different normative evidence.) Ignorance, on his view, does not excuse the wrongs or bad states of affairs the agent did not foresee she would bring about. Ignorance obviates the need to justify conduct that brings about bad outcomes.

In Chapter 1, Zimmerman makes heavy use of variants of Regan's cases to argue in favor of the Prospective View of moral obligation and against the Objective View:

PV: An agent ought to perform an act if and only if it is the prospectively best option that she has.

OV: An agent ought to perform an act if and only if it is the best option that she has.

In Chapter 2, he provides a prospectivist account of *prima facie* obligation and moral rights. On his view, others violate our rights by imposing risks upon us but not by harming us, *per se*. An apparent consequence of this view is that agents who harm others do not owe compensation or reparation unless these agents' actions were risky from the point of view of these agents. The agent's

---

<sup>1</sup> Zimmerman (1996: 13). For defenses of the objectivist view of 'ought', see Thomson (1986: 179) and Moore (1903).

<sup>2</sup> Zimmerman (2008: x).

<sup>3</sup> See Regan (1980). The examples that Zimmerman discusses are often taken from Jackson's (1991) discussion and modified for his purposes.

perspective is privileged, not the victim's. In Chapter 3, he extends the account to explain why we should prefer Prospective Possibilism to Actualism. The actualist thinks that an agent ought to perform an act if its performance *would* be better than what would happen if the agent chose to do something else instead. The possibilist thinks that we have to consider what *could* happen if the agent acted and whether things could have been better if the agent chose to act otherwise. An actualist will say that facts about what an agent will decide to do later can have some bearing on present obligation. Possibilists deny this. Finally, in Chapter 4, Zimmerman defends the view that an agent cannot be blamed for acting from normative or non-normative ignorance unless the agent is culpably ignorant. He then argues that we will rarely be culpable for our ignorance as all culpable ignorance must derive from an act the agent performed in the belief that the act was wrong.

While I find many of Zimmerman's arguments persuasive, I do worry that the moral 'ought' is messier than Zimmerman's account would have us believe. In what follows, I shall try to identify where I think Zimmerman's account runs into trouble.

### 1. THE PROSPECTIVE VIEW

Let's start with a case that has the potential to cause trouble for the Objective View:

#### *Case 1*

All the evidence at your disposal indicates that giving the patient drug B would cure him partially (+40), giving him nothing would leave him permanently incurable (0), giving him drug C would cure him completely (+50), and giving him drug A would kill him (-100). What should you do?

An objectivist would likely say that you ought to give the patient drug C. Suppose, however, that you give the patient drug C and it kills him. It turns out that your evidence was misleading and it was drug A that would have cured him. What can the objectivist say in response? Does the fact that most of us would say that we ought to give the patient drug C if put into a situation like this show that OV is mistaken? It might, but the objectivist has a reply at the ready. She would say that mistakes were made and that you really should not have given the patient drug C. The mistake was excusable, however, because the evidence indicated your action would have been for the best. You did *try* to do the best and *try* to do what the objectivist says you ought to do. Who could fault you for trying your best?

A slight variant on the above case seems to show that the objectivist's response to *Case 1* is incomplete at best:

#### *Case 2*

All the evidence at your disposal indicates (in keeping with the facts) that giving the patient drug B would cure him partially (+40). The evidence does not tell us whether it is drug A or drug C that will kill the patient (-100). The evidence does not tell you whether it is drug A or drug C that will cure the patient completely (+50). You do know, however, that one of these drugs would kill and one would cure. What should you do?

Zimmerman thinks that everyone (including the ghost of G.E. Moore) will say that you should give drug B knowing full well that this is not the objectively best option. If, as he suggests, the conscientious moral agent will never deliberately do what the agent believes to be overall morally wrong, the objectivist view is in serious trouble. Objectivists who think they should give drug B are acknowledging that *they think* the Objective View is mistaken. It is not as if the objectivist can

plead ignorance in this case or say that in giving drug B, the agent is *trying* to do what is objectively best. The objectivist believes that giving the one drug certain not to bring about the best results is drug B.

An objectivist who wants to blunt the force of Zimmerman's objection might caution us against drawing any general lessons about ignorance and moral obligation from cases like *Case 2*. An incorrigible objectivist might say that there is an important difference between *Case 2* and *Case 1*. In *Case 2*, the objectivist could say, the agent knows that she cannot know what she ought to do and is forced to make a choice knowing also that if she tries to do the best there is a very significant chance of bringing about a very significant harm. In *Case 1*, this isn't the case. Does this difference matter? Maybe the agent says what she says in response to *Case 2* because she's trying to minimize risk and *not* because she believes she is really obliged to give drug B. Zimmerman considers this sort of response and notes (rightly) that we cannot say that she's trying to minimize risk of wrongdoing. (Giving drug B maximizes the chance of wrongdoing.) Perhaps the objectivist should say that the agent is trying to minimize the risk of *harm*. She might be thinking that since she does not know what she (really) ought to do, the next best thing she could do is minimize the risk by aiming for the next best thing. If that's right, perhaps *Case 2* is a special case where 'ought' is being used in some special way and the agent who deliberately gives drug B saying this is what she 'ought' to do is not the demise of the Objective View.

Does this force the objectivist to say that 'ought' is ambiguous? That is one way they could try to go, but they might not have to. Perhaps in judging that she should give drug B she's really thinking something like this: *if I know I don't know whether it is drug A or C that is best and know that guessing could be disastrous, I ought to give drug B*. The objectivist might say that the judgment that drug B should be given in *Case 2* is a judgment about conditional obligation.<sup>4</sup> Consider an example. An advisor might know that the best thing (objectively and prospectively) for an advisee to do is something the advisor learns the advisee will not do. The advisor might then say she ought to do something else, the next best thing. Does that mean that this is not a conscientious advisor? No, the morally conscientious advisor can deliberately tell an advisee to do what the advisor believes is overall morally wrong *when there's good reason to do that*. To make this somewhat concrete, I have an (imaginary) advisee who wants to do graduate work in philosophy. I think this isn't a very good idea. I think she should go into law.<sup>5</sup> While I think we both know that it would be best for her to do that, I think I'm still a perfectly conscientious advisor if I say that she should go do graduate work at a highly ranked program in philosophy. I know that she will not apply to law school and so will only choose to do something suboptimal, so if she is going to do that, I should advise her to do the least suboptimal thing she's willing to do. There's a good reason to give the advice I'm going to give. The advice helps her identify the second best option. While I know she could do better, I know that she simply won't do what's best. If the objectivist can find a good reason for the agent to advise *herself* to do other than what she believes is overall best, perhaps the objectivist can agree that the conscientious agent would say what Zimmerman says she'd say and still remain an objectivist about overall unconditional obligation. So, can the objectivist say something similar about *Case 2*?

---

<sup>4</sup> Zimmerman (2008: 61) explains why he thinks the distinction between conditional and unconditional obligation will not be useful for the objectivist.

<sup>5</sup> The reasons for her to prefer law school to graduate work in philosophy need not concern us here, but it should not be hard to imagine what they might be (e.g., it is hard to find work in philosophy, philosophers face the two body problem more often than lawyers do, graduate school is a place for fanatics, people with families have a harder time doing seven years of graduate work than three, etc...).

Can she say that what's really going on here is that we'd say that what the agent (really) ought to do is give either drug A or C but since we don't know which drug and we know we shouldn't give both, it would be very risky to try to do what is best for that reason we deliberately decide to go for an option that is safe to as to avoid making a bad situation worse? What we (really) believe is *not* that we have an unconditional obligation to give drug B, but a conditional obligation, one that is conditional on our ignorance.

It is hard to know what a principled objectivist view would look like if it tried to accommodate the thought that a conscientious agent in a case like *Case 2* could properly/correctly/truthfully say that she should give drug B. At the very least, the objectivist would owe an account of conditional obligation as the account Zimmerman provides would not be suitable for the objectivist. Perhaps the objectivist can show that the problem she faces is a problem that arises for the prospectivist as well. This would force the prospectivist to modify her own view or rethink the assumptions that figure in the argument against the objectivist view. The problem for the objectivist might be put like this. Suppose that according to theory<sub>x</sub>, an agent ought<sub>x</sub> to do what is best<sub>x</sub>. Suppose an agent is faced with three options (A, B, and C) and the agent knows she ought<sub>x</sub> to do A or C because she knows that A-ing or C-ing is best<sub>x</sub> but doesn't know which. An advocate of theory<sub>x</sub> cannot then say that the agent in such a predicament ought<sub>y</sub> to do B provided that doing B is best<sub>y</sub> if what an agent ought<sub>x</sub> to do and ought<sub>y</sub> to do are distinct because there's a difference between the best<sub>x</sub> and the best<sub>y</sub>. So, if you construct a case where the conscientious agent judges, 'I should do what is best<sub>y</sub>' in the knowledge that it's not best<sub>x</sub> and that seems to us like the thing to say, it's not the case that the agent *really* ought to do what theory<sub>x</sub> says (except on those occasions where the best<sub>x</sub> is the best<sub>y</sub>). For someone like Moore, what you *really* ought to do is what is best<sub>objectively</sub> and we see what the objection does to his view. For Zimmerman, what you *really* ought to do is what is best<sub>prospectively</sub> and here's a case that I think is troubling:

*Case 3*

Your patient is ill. If nothing is done, she will feel very ill for a few days (-100). You know that if you give her drug B, she will feel much better (+40). Your old probability and statistics professor appears. He has with him two drugs, drug A and drug C. One drug cures quickly (+80). One drug kills slowly (-200). He hands you both drugs. You ask which drug cures. He says that he will tell you if you can tell him the probability that a defective widget drawn at random came from the factory in Betaville. He then gives you all the facts you would need to determine that. If you tell him the right answer, he promises to tell you which drug cures. If you tell him the wrong answer, he will *say* that one of the drug cures but he will flip a mental coin to determine whether to tell you the truth. (He then tells you that 40% of the widgets come from a supplier in Alphatown and 60% from a supplier in Betaville. 3% of the widgets from Alphatown are defective and 6% of the widgets from Betaville are defective.) You give it your best shot, you tell him that the probability is 75%, and he tells you that it is drug A that cures. What should you do?<sup>6</sup>

---

<sup>6</sup> Incidentally, I did not do the math to determine if 75% is the correct answer. This is left as an exercise to the reader to feel the force of the relevant intuition.

On one way of understanding what is prospectively best, you know that the prospectively best thing for you to do is give the patient one of the drugs that your professor gave you, the drug that gives the complete cure. If you are anything like me (i.e., you can pass your math classes and that's the best thing your former math professors could say about you), it would be crazy to give your patient the drug that is best<sub>prospectively</sub>. The chances that I would not correctly identify the option that is best<sub>prospectively</sub> and kill the patient are too great and I know I have the option that is second best<sub>prospectively</sub> and second best<sub>objectively</sub>. I think I could deliberately choose and say correctly/truthfully/accurately 'I should give the patient drug B'. There are passages where Zimmerman says that it is not easy to know which option is prospectively best because it is not easy to know which option will maximize expectable value.<sup>7</sup> If we take him at his word, I worry that our intuitions about *Case 3* will cause essentially the same problem for the Prospective View that our intuitions about *Case 2* are supposed to cause for the Objective View.

It is not entirely clear what Zimmerman thinks our evidence consists of.<sup>8</sup> He might say that in this case you should fold into your evidence further evidence about how likely it is that you will respond correctly to your evidence in this situation. Does this help the mathematically challenged identify which option is best<sub>prospectively</sub>? Some of us do not have much evidence about how likely it is that we respond correctly to evidence in cases like this, but we might have varying degrees of confidence in our ability to correctly run the numbers. If these varying levels of confidence are part of what determines which option is best<sub>prospectively</sub>, do we end up saying that the agent's obligations will depend upon whether they are confident in answering the professor's question? Because we can easily imagine people who are overly confident or insufficiently confident in their abilities, I worry that pursuing this line of response leads to a view that will say that overly confident agents should do what they shouldn't and insufficiently confident agents shouldn't do what they should. It seems we have a recipe for constructing cases that cause trouble for the Prospective View. Speaking just for myself, the intuition that Zimmerman uses to convince me to abandon the Objective View looks a lot like the intuition I have about *Case 3*. If that intuition is not a problem for the Prospective View because the view does not imply that the agent ought to give the patient the drug that provides the complete cure, I worry that our inability to formulate counterexamples to the view is due to the fact that we cannot say what an agent's evidence consists of and so cannot work out the implications of the view to test it by appeal to intuitions about cases. While I have to admit that it's hard to see what a principled and intuitively plausible objectivist view would look like, I think we have some reason to think the issue is just a bit messier than the Prospective View suggests.

Let's grant that intuitions about *Case 2* make it tempting to embrace something like the prospectivist view. Let me briefly note two more concerns. First, while there clearly are intuitions that favor the prospectivist view, there are also intuitions that seem to favor the objectivist view. Suppose in *Case 2*, an advisor who knows more than you happens to know which drug will cure the patient and so says, 'You should give drug A'.<sup>9</sup> It seems she speaks correctly/properly/truthfully, etc... (Of course, the agent in *Case 2* prior to this could say to herself, 'I should give drug B' and it seems she speaks correctly/properly/truthfully, etc...) Some say that the advisor's remarks are proper because those remarks, if followed, would lead to a better outcome. That's true, but that's not the point. Even if the advisor knows that her advice will go unheeded (e.g., she knows that the advisee will likely disregard her advice because there are too

---

<sup>7</sup> Zimmerman (2008: 44).

<sup>8</sup> Zimmerman (2008: 35).

<sup>9</sup> See Thomson (1986: 187).

many bad advisors in the neighborhood), she knows that she speaks truthfully/properly/correctly in saying that the agent 'ought' to do other than what the Prospective View says. Zimmerman rejects the idea that 'ought' is ambiguous. He probably should. An ambiguity thesis would not accommodate the relevant intuitions. It seems to me and would seem to the advisor that she is *disagreeing* with the agent when she says the agent ought to give drug A.<sup>10</sup> This is hard to make sense of on the hypothesis that 'ought' just means different things coming off of the lips of our two speakers. Shouldn't we have some account of what's going on here? We need a story that either accommodates or explains away the intuitions that favor the objectivist view, but we haven't been given one.

Second, the intuitions that have the greatest force do not motivate Zimmerman's version of the prospective view because his view treats normative and non-normative uncertainty on par.<sup>11</sup> Suppose an agent is trying to decide which of two vaccines will be administered to his daughter. Both are equally effective as a precautionary measure against cervical cancer, let's say, but one has the added bonus of being an effective precautionary measure against HPV. The father takes this to be a con. He thinks this diminishes the health risks of premarital sex and for this reason the father judges that he ought to have the first vaccine administered. I don't know really how to describe the agent's evidence in cases like this. I know I'm supposed to take account of both normative and non-normative uncertainty and the sad fact is that this agent is not certain that he ought to protect his daughter from HPV because he is certain she should be at greater risk for STDs to deter her from premarital sex.<sup>12</sup> I cannot say that I know that his action does not maximize expectable value given his very defective perspective on things, but don't I know that he fails to do what he should as a father? The worry is this. As the agent's own views about the right and the good start to deviate further and further from the facts about what is right and good, the view that identifies the agent's obligations with the thing that maximizes expectable value given *their* 'evidence', won't the prospectivist end up sanctioning wrongdoing? I don't think our intuitions about cases like *Case 2* do much to motivate this sort of view even if some of our intuitions suggest that there is some role for normative ignorance to play in determining what our obligations are.

## 2. RIGHTS AND RISK

In keeping with his prospectivist perspective, Zimmerman rejects the Harm Thesis:

HT: We have moral rights against others that they not cause us harm.

In its place, he defends the Risk Thesis:

RT: We have moral rights against others that they not impose risks of harms on us.

The argument against HT and in support of RT is straightforward. First, according to the Prospective View, what an agent ought to do is whatever happens to be prospectively best. According to the Correlativity Thesis:

---

<sup>10</sup> See Kolodny and MacFarlane (forthcoming) for discussion.

<sup>11</sup> Zimmerman (2008: 38).

<sup>12</sup> Yes, I'm sad to say, this is an actual view held by actual people who have actual children. They cite the view when they try to explain why they will not allow their children to be given immunizations that they do believe reduces the risk of cervical cancer. I do realize that there is the very real risk of contracting an STD even if you never have premarital sex. My (imaginary) father is not very bright and not very good. So far as I know, the only vaccines that protect against cervical cancer also protect against HPV and so the actual fathers that my imaginary father is based on are people who deny their daughters any protection against cervical cancer on the grounds that it makes sex safer.

CT: One person, Q, has a moral right against another person, P, that P perform some act, A, if and only if P has an obligation to Q to perform A.

To determine whether someone has had her rights violated, we need to know something about the epistemic situation of the agent who harms rather than the subject harmed. If the agent did what was prospectively best but their actions caused harm, Zimmerman insists that this action violates no rights. It's unfortunate when this happens, but that's a different thing entirely.

I think a good case can be made for HT. Suppose Mustard makes dinner for Plum to welcome her to the neighborhood. His dish contains poisoned mushrooms, but we might fill out the details of the story in such a way that Mustard's behavior was prospectively best by imagining that he had good evidence that the mushrooms he picked were safe. Because she ate his dish, Plum is now violently ill. Another of Mustard's neighbors, White, is violently ill because of food poisoning. (He is ill for reasons that have nothing to do with Mustard.) Mustard has on hand the stuff that would fix Plum and White up but only enough to help one. It seems to me that Mustard has a more stringent duty to assist Plum than White. This difference in the comparative strength in duties suggests that Mustard's duty to Plum is no mere duty of beneficence. (If it were and we assumed that Plum and White were in equally bad shape, I would think that if the drug would help them equally there would not be a difference in the stringency of duties if both were duties of beneficence.) I don't think the prospective view could make sense of the idea that Mustard's duty is a reparative duty because there is nothing that Mustard did on that view that is overall wrong or *prima facie* wrong.

While I think these intuitions suggest that Plum has a right to Mustard's assistance and so cause trouble for a view that denies HT while holding onto CT, some of Zimmerman's remarks suggest that he doesn't think that Mustard has to compensate or make reparations. Against the claim that someone is owed compensation by those who harm them, Zimmerman says that (i) this leaves some needy parties (e.g., White) "out in the cold" even if this party is just as deserving of compensation, (ii) the party that harmed may have been just as innocent as the party harmed, and (iii) there might be some fourth party who is just as much at fault as the party that causes the injury that is just as deserving to be made to make amends who we know should not be made to do so.<sup>13</sup> This is not, to my mind, a convincing line of response. I don't think we can determine what someone's obligations are by determining whether we think there is independent reason to think that they deserve to be under these obligations or made to live up to those obligations. Against (ii) or (iii), no one deserves to be under a duty of beneficence, but we are for that often duty bound to assist others at an expense to ourselves and when we are perfectly innocent in terms of what brought it about that they need our assistance. Against (i), I think we cannot rest too much weight on this point. Suppose Mustard had tried to poison Plum and succeeded in so doing. If White and Plum are equally faultless in finding themselves poisoned, surely they are equally deserving of assistance, but nobody would say that Mustard's obligations to Plum are for that reason not stronger than the duties he has to those he has not tried to kill. If (i) were applied consistently, I think it would essentially prevent us from saying that victims are owed compensation by those who put them at risk of harm for no good reason just as surely as it would prevent us from saying that victims are owed compensation for being harmed with no overriding reason to have done that.

There is some motivation for rejecting HT and putting RT in its place. Taken in combination, HT and CT force us to adopt some sort of Objective View that Zimmerman thinks he's dispensed with in Chapter 1. If we rely on intuitions about compensation and reparation,

---

<sup>13</sup> Zimmerman (2008: 84).

however, I think a pretty good case can be made against a view that denies HT but accepts CT. I did not find the response to this kind of argument altogether convincing. (Perhaps I can take some consolation in the thought that given the high level of credence I place in HT, I will owe compensation to those I know I have harmed even if others get off of the hook.)

### 3. PROSPECTIVE POSSIBILISM

In his earlier work, Zimmerman built a powerful case against Actualism:

AO: An agent ought to perform an act if and only if it is an option such that what *would* happen if the agent performed it is better than what *would* happen if she did not perform it.

He still defends Possibilism:

PO: An agent ought to perform an act if and only if it is an option such that what *could* happen if the agent performed it is better than what *could* happen if he did not perform it.

One of the problems with AO is that it implies that you can escape certain obligations now if you are the sort of person who would do wrong in the future. Some of us would live up to our obligations to get a referee's report in on a paper we agree to referee and so are obliged to accept an invitation to referee if we are the best person for the job. The actualist view says that this sort of professional obligation is not one our irresponsible colleagues are under if they are the sort of person who would not complete the review if they accepted the invitation to give one (provided that it would be better to decline and let a more responsible colleague shoulder the burden). As it stands, PO is not stated in prospectivist terms. Modifications are needed. Prospective Possibilism (after much refinement) comes to this:

PP: P ought at T to do A at T\* if and only if:

- (1) P can at T do A at T\*,
- (2) P can at T refrain from doing A at T\*, and
- (3) for every maximal course of action, C, that P can at T perform and which excludes P's doing A at T\*, there is some maximal course of action, C\*, such that:
  - (a) P can at T perform C\*,
  - (b) C\* includes P's doing A at T\*, and
  - (c) C\*'s core is prospectively better, for P at T, than C's.<sup>14</sup>

I'm not certain that PP accommodates all of our intuitions. It takes a complicated case to cause trouble for PP. I hope the reader will forgive me for *Case 4*:

#### *Case 4*

You have a sick patient. You can take her to one of two miracle workers. Miracle Max in Alphatown could cure her completely (+36). Miracle Minnie in Betaville could cure her partially if you travel by the southern route (+8) and cure her partially if you travel by the northern route (+6). You can travel the northern route and reach Alphatown, Betaville, or Gamma City. You can

---

<sup>14</sup> I've introduced a healthy dose of jargon. Let's say that a course of action, C, will have a core that consists of all the attempts that the agent can make (intentionally) in carrying out that course of action. A course of action, C, is a maximal course of action performable by the agent just in case there is no other course of action that is performable by the agent that includes C.

travel the southern route to Betaville. You are advised to stay away from Gamma City. There your patient will only find certain death (-300). You know that hundreds of thousands have travelled along the northern route to find Max. You know that each traveler has decided to head down to Betaville rather than travel on to Alphatown or Gamma City. You know that each traveler that has headed north has brought about the third best outcome (characterized objectively). You know the reason for this. The travelers who head north forget whether Max is in Alphaville or Gamma City. Knowing the risks of guessing, they head to Betaville for a partial cure (+6). You think that the probability that you will forget and head to Betaville if you head north is 1. Expert opinion is evenly divided as to why travelers forget. Half of the experts think that the northern route is lined with temptation. The travelers have all given in to some moral weakness (e.g., a weakness for poppies, for drink, etc...) and the result is that they lose the evidence they need to get to Max. Half of the experts think that the northern route is lined with nogoodniks in Minnie's employ. They use their considerable powers to muddle the minds of the travelers who forget where Max is. Knowing the risks of guessing, the theory is that the travelers decide to head to Minnie in Betaville from the north. You dutifully divide up your credence. You think the probability that you will forget due to some moral lapse is .5 and the probability that you will forget due to something other than a moral lapse is .5. You know where you will end up even if you do not know why. Should you head north or south?<sup>15</sup>

I think it's clear that you should head south rather than north, but I think this isn't what PP says. There's a 50% chance that if you head north you could see to it that you do not suffer any moral lapse and retain the information you need to know Max's location (+36). (You will not do that, of course, but you could.) There's a 50% chance that if you head north, you would lose the information that would get you to Max even if you do not suffer any moral lapse. If that happens, you know that you could get to Minnie and would go to Minnie (+6). If you head to Minnie by the southern route, you can get Minnie to partially cure your patient (+8). It seems on PP, the expected value of heading north (+21) is greater than the expected value of heading south (+8). Of course, you know that it is nearly certain that if you head north, the outcome you will bring about is that you will head to Minnie in a costly way (+6) rather than a cost-free way (+8).

---

<sup>15</sup> What matters for the possibilist is what you *could* do if you head north. To keep the math simple, I have assumed that the probability that you *will* lose the evidence you need to know Max's location is 1. I think this is compatible with the further claim that you *could* head north and not suffer from any moral lapse that would result in the loss of evidence. If it is not, you might change the case as follows. Of the hundreds of thousands that headed north, one pillar of virtue managed to remember where Max was. Experts are divided as to whether it was this man's moral virtue or just a failure on the part of Minnie's minions. The probability that someone could head north and remember Max's location, we might say, is something vanishingly small but greater than 0. You should still head south.

(There's a 50% chance that you will do this of your own free will and a 50% chance that you will do this because your autonomy has been undermined. The problem is that Prospective Possibilism does not take account of the probability that you will fail to bring about an outcome or attempt to bring about that outcome due to a moral lapse so as to avoid collapsing into a probabilized version of Actualism.) I think there's a good case for the Prospective View and a good case for Possibilism, but I think Prospective Possibilism needs refinement. I don't see a good way to do that.<sup>16</sup>

#### 4. IGNORANCE AND MORAL RESPONSIBILITY

In the book's final chapter, Zimmerman defends what is (to my mind) a radically subjectivist account of moral responsibility. On his view, an agent cannot be blamed for acting wrongly unless she acts in the belief that her action is wrong or she fails to believe that her act is wrong because of something that she is directly responsible for that she did in the belief that *that* was wrong. The argument for this view is essentially this. Suppose the agent's action was wrongful. If the agent did not believe that the action was wrongful, the agent was ignorant of the fact that it was wrongful and no one can be culpable for acting from ignorance unless that ignorance is itself something for which the agent is culpable. But, whatever it is that we trace culpable ignorance to, it had better be something that the agent is culpable for, which requires that the agent did that in the belief that it was wrongful.

Zimmerman does believe that lack of belief typically excuses all manner of behavior. He does believe that the presence of belief can render the agent culpable for action, say, when the agent acts in the belief that the action is wrong. Suppose I come to believe that I have quite extensive obligations to the poor. I make extensive sacrifices to live up to my own very demanding moral standards. Every so often, I indulge in something minor (e.g., an ice cream cone) and when I do so, I do so in the belief that I act wrongly I know I could have better used my resources in the service of the poor. I can be blamed for my minor indulgences to a greater degree than an unreflective murderous tyrant can be blamed simply because the thought that the mass killing of innocents is wrong is a thought that never crossed his mind. (Of course, you can say that the thought *should* have crossed his mind, but the issue is whether he is culpable for the fact that it did not and so culpable for the fact that a thought that should have crossed his mind didn't.) Zimmerman reminds us that we can say that someone who perpetuates evil without believing that this is what they are doing can be said to be 'reprehensible', but I have to admit that I don't have a firm grip on the idea that we can reasonably say that an agent is morally reprehensible for doing things we ourselves know to be things the agent is in no way morally responsible for.

---

<sup>16</sup> We can generate interesting test cases by looking at cases where an agent knows that her values will change through the course of some action and so has to decide how to accommodate future changes in normative evidence. (We can imagine cases inspired by Pascal's observation that those who manipulate themselves to become believers end up with values at the end of the process quite different from those that would lead them to believe that it is a good idea to manipulate themselves in order to receive a reward for their efforts.) If I have good evidence that my normative evidence will change over a longer course of action, how should I plan for that now knowing that what might be prospectively best at the beginning of my journey might not be in the middle or near the end? In *Case 4*, I think there's some reason to think that the agent's initial body of evidence is what determines what the agent should do from beginning to end because there's an obvious way in which the initial evidence is just better than the evidence the agent has if she heads up the northern route. It's hard to know how to evaluate different bodies of normative evidence if we do not have some independent way of determining which body of evidence is 'better'.

In acting ignorantly, an agent might fail to believe *of* some feature she's aware of that it is a wrong-making feature or she might fail to believe *that* some wrong-making feature is a feature of her action or the circumstance in which she acts because she's unaware of the feature. I don't understand why Zimmerman thinks this is a distinction without a difference when it comes to ascriptions of moral responsibility. It is true that we do say things that suggest that we sometimes think that an agent who is non-culpably ignorant of the fact that she's engaged in wrongdoing is blameless for having done what she did, but it's hardly obvious that the agent who is aware of the wrong-making features without being aware of them *as* wrong-making features gets off the hook as easily as the agent who has not been made aware of all of the relevant non-normative facts. The agent who acts in awareness of the reasons that make a decisive case against acting but acts anyway is the agent who is willing to act against reasons that are, well, reasons. The agent's deeds show that she has the wrong values. That's why it is hard to excuse the agent's conduct. I'm sure this needs refinement. It's too crude as stated. However, the view that says there is an asymmetry here in the way that we take account of normative and non-normative ignorance in ascriptions of responsibility is a common one. That someone identifies with the wrong values may well explain her disposition to act wrongfully and her disposition to do this with a clear conscience, but it's hard to see how it explains why she acted blamelessly. In my view, I don't think Zimmerman has said enough to address this kind of view in this discussion.<sup>17</sup>

#### REFERENCES

- Jackson, F. 1991. Decision-Theoretic Consequentialism and the Nearest and Dearest Objection. *Ethics*: 101: 461-82.
- MacFarlane, J. and N. Kolodny. Forthcoming. Ifs and Oughts. *Journal of Philosophy*.
- Moore, G.E. 1903. *Principia Ethica*. New York: Cambridge University Press.
- Regan, D. 1980. *Utilitarianism and Co-operation*. New York: Oxford University Press.
- Thomson, J. 1986. *Rights, Restitution, and Risk*. Cambridge, MA: Harvard University Press.
- Zimmerman, M. 1996. *The Concept of Moral Obligation*. New York: Cambridge University Press.
- . 2008. *Living with Uncertainty: The Moral Significance of Ignorance*. New York: Cambridge University Press.

---

<sup>17</sup> I want to thank Mike Almeida and an anonymous reader for this journal for helpful feedback.