Michael J. Zimmerman, *Ignorance and Moral Obligation*, Oxford University Press, 2014, pp. 149, £ 30.00, ISBN 9780199688852

Simone Grigoletto, Università degli Studi di Padova

How does ignorance of the circumstances and about ourselves affect our moral obligations? What if the facts that we happen to consider relevant are not enough to answer the question of which is *actually* the best option? Zimmerman's *Ignorance and Moral Obligation* combines previously-published and original work in an effort to answer such questions. The book contains a series of arguments, objections and replies that aim to shed light on which of the different views on moral obligation is preferable.

In chapter 1, Zimmerman distinguishes between three different possible accounts of a moral obligation: the first is the *Objective View* (*OV*), and that is, "*one ought to perform an act if and only if it is in fact the best option that one has*" (p.2); second comes the *Subjective View* (*SV*), namely, "*one ought to perform an act if and only if one believes that it is the best option one has*" (p.7); and third there is the *Prospective View* (*PV*), or "*one ought to perform an act if and only if it is probably the best option one has*" (p.8). In short, the three views express how to identify an agent's moral obligation, once we have chosen what is morally at stake. The outcome of the three views coincides whenever the agent – based on the available *evidence - knows* for sure what it is *actually* best (and therefore *believes* in the choice). Zimmerman calls this a *happy case* (p.22). But what happens if we face a less fortunate (and probably more common) situation where there is a shortage of *evidence* on what is really the best option to pursue? This is what Zimmerman calls *Question 1*: "*What ought one to do when one doesn't know which of one's options is best?*" (p.10).

Chapter 2 tries to deny the effectiveness of the *OV* and *SV*, providing reasons for justifying the *PV*. At the start of this chapter, Zimmerman deals with a major objection to the *SV* that he calls *Constraint #1* (p.27). According to this objection, it is perfectly reasonable for an agent to have a misguided conscience (or, as we might say, for an agent to ignore the evidence), so we cannot agree with the view that what one *ought* to do is what one *believes* to be morally best. Agents might hold

false beliefs, so their moral obligations cannot be grounded on them. In this context at least, moral obligations refer to the *actual value* of our available options and *Constraint #1* warns us that it is perfectly reasonable for even the conscientious agent not to have access to this value (p.28).

The second chapter then goes on to reject the *OV*. Zimmermann provides many cases and examples throughout the book, but certainly one of the most recurrent is *Case 2* (that he draws from Frank Jackson). This case is used to show why the *PV* has some advantages over the *OV*. The scene depicts Jill struggling with a decision about how to cure John of a potentially deadly skin complaint. This is the first version of the case, which I quote in full due to its relevance: "*All the evidence at Jill's disposal indicates (in keeping with the facts) that giving John drug B would cure him partially and giving him no drug would render him permanently incurable, but (despite the facts) it leaves it completely open whether it is giving him drug A or giving him drug C that would cure him completely and whether it is giving him drug A or giving him drug C that would kill him. Heeding the available evidence, Jill gives John drug B, thereby providing him with partial cure*" (p.30). The *OV* implies that what she ought to do is either option A or C, regardless of the fact that she does not have enough evidence to tell which drug will kill her patient and which one will cure him completely. The most reasonable thing to do here was to give John drug B, despite the fact that, from the *OV*, this is not what Jill ought to have done (because it was not the best option). Zimmerman proposes a *refined* version of the *PV*: "*One ought to perform an act if and only if it is one's prospectively best option*" (p.32). By *prospectively best*, he means the option that would be more likely to improve John's health and that would avoid what might turn out to have a negative value, i.e. John's death (either option A or option C, based on Jill's evidence). This is an example of what the author calls *Constraint #2*: if one acts morally conscientiously, then one does not do what one believes to be morally wrong (p.32). Since the agent who acts morally conscientiously would never pick an option other than B in case 2, the *OV* (indicating that this is not what ought to have been done) has to be rejected. Unfortunately, one of the few weaknesses in Zimmerman's argument is that he does not spend much time on telling us *why* a morally conscientious agent would not risk choosing an option other than B. Probably, some

mention of the many studies underscoring human beings' distinctive risk aversion in decision-making situations might have helped to support this line of argumentation.

More needs to be said about the third and last view. The *prospectively best* option could be interpreted as the agent's *best bet* on the available options. This is the one that maximizes the *projected value* of an act, expressed "*as a function of the probabilities of the act's possible outcomes and the* probable, *rather than actual, values of these outcomes*" (p.36). We can therefore leave aside the *actual* value of an option, since the evidence available to the agent might not be enough to grasp it. Having said that, the *PV* aims roughly at the prospectively best option, or the *best bet*, i.e. the one that maximizes the *projected* value.

The remainder of the second chapter and the whole of the third address twelve objections to the *PV*. The author tries to respond to each objection in further defense of the *PV*. For reasons of space, I will mention only the two that I consider the most pressing against the author's position.

Objection 6 points out that the *PV* might imply some horrific outcome (pp.60-61). Consider a case where Jill, based on the evidence at her disposal, believes that drug A would kill John and drug C would cure him completely. In virtue of this, she gives him drug C. It turns out that drug C kills John and that drug A would have completely cured him. According to the *PV*, she ought to have given him drug C but this seems unacceptable. In response to this objection, Zimmerman starts by emphasizing that the *SV* would have prompted the same behavior. He adds that the *OV* can generate counterintuitive outcomes too. In fact, as we have seen in *Case 2*, according to the *OV*, Jill should have given John one drug despite the evidence available to her suggesting another. Zimmerman concludes from this that, although all three views may have counterintuitive outcomes, this does not justify rejecting all three – given that a view that escapes such criticism does not exist. After all, he says, his rejection of both the *SV* and the *OV* are based on different reasons (p.61).

Objection 11 states that, in situations such as in *Case 2*, the effectiveness of the *PV* is undermined when agents ask someone for advice on what they should do and discover that their original choice was indeed misguided. In the case in point, taking the prospective view, Jill is convinced that she ought to

give John drug B, even though she knows it is a suboptimal option. When she consults her colleague Jack, she learns that the best option is drug A. We might therefore conclude that the *PV* was originally misleading in pointing to the agent's moral obligation (p.82). Zimmerman replies that, when Jill asks Jack for advice, she does not ask for a *prediction* of what her proper moral obligation would be; she is asking instead for his cooperation in identifying her moral obligation (p.86). This answer seems reasonable. After all, the helper's advice, when appropriately given to the agent, plays a relevant part as further evidence. It suggests that choosing drug A is the *best bet*. In the scenario depicted, Jack's knowledge *grounds* Jill's moral obligation, since it has become Jill's knowledge as well.

In the fourth chapter, the author shifts his attention to the problem of action-guidance. Most of the normative theories in ethics, he claims, have a feature in common, and that is an *objective* approach to what matters morally. Kantianism, virtue ethics, and different versions of utilitarianism all entail that one ought to choose the *best* option one has, based on what matters morally. The whole book aims to show that, most of the time, ignorance plays a relevant part in decision-making situations, i.e. the agent might not know which option is actually best.

Among all the normative theories, the chapter focuses mainly on the downsides of a simplified version of act-utilitarianism. Zimmerman recalls a common objection to this theory: it often happens that the agent does not know which of the available options would maximize utility. The agent is also rarely in a position to calculate this utility reliably. In the past, Utilitarians would have responded to this objection that, when this is the case, we can resort to some commonsensical rules that generally point towards the satisfaction of the principle of utility. Let us call it *Guide #1*. This acts as a supplement to the general principle on which the normative theory is grounded. Now, in the spirit of the *PV*, Zimmermann claims that it is hard to apply *Guide #1* successfully to the principle of utility because this might involve us *knowing* exactly what the principle considers *best*. The author's move is to introduce a weaker sense of the term "apply", the *enactment* sense (p.103), according to which the agent does not *know* what the principle prescribes, but does *justifiably believe* that the chosen option is prescribed by the principle. This refinement takes advantage of the fact that the *degree of justification* of one's belief need not to be as strong as

the one required for knowledge (which is hard to achieve) (p.104). This chapter tries to show that it is particularly in cases of action-guidance that the *PV* reveals all its advantages over the *OV* by endorsing a weaker (but less problematic) sense of application of the general principle.

The last chapter is concerned with a possible objection to the *PV* that arises from the so-called *Correlativity Thesis* (CT) of moral rights. This thesis roughly states that it is someone's right to have another perform a certain act if and only if the latter person has a moral obligation to the former to perform said act (p.116). Moral rights and moral obligations are therefore closely related. In fact, consistently with the *PV* – according to which the agent's moral obligations depend on the evidence available to her, moral rights happen to follow the same path (p.126). Zimmerman not only accepts the *CT*, but also claims that our rights depend on other people's *ability* to satisfy them and, most of all, on other people's *epistemic circumstances*, which ground their obligations (and with them, our corresponding *rights*). This conclusion might seem disheartening, but Zimmerman tries to sweeten the pill by drawing a clear distinction between *desert* and *rights* (pp.123-124). Though we might not have the *right* to be treated justly, that does not mean to say that we do not *deserve* to be treated justly (p.127).

Generally speaking, despite what the author writes in his Preface, it is very helpful (if not necessary) to have some background knowledge of the author's previous publications in order to understand some passages of this book. Overall, we acknowledge the merit of this work as an attempt to underscore how *ignorance* might play a major part in decision-making situations. An understanding of moral obligation that takes this into account is therefore the only way to avoid making our moral systems too demanding, or even scarcely feasible.

Bibliography

Michael J. Zimmerman, *The Concept of Moral Obligation*, Cambridge University Press, 1996.
Michael J. Zimmerman, *Living with Uncertainty*, Cambridge University Press, 2008.