

Conclusion

This book has covered a lot of ground. In the first chapter, we gave the case for taking moral uncertainty seriously, and for thinking that there are non-trivial answers to our central question, ‘Given that we are morally uncertain, how ought we to act in light of that uncertainty?’

In Chapters 2–6, we developed a general account that provided our answer to that question. We defended an information-sensitive account: the correct rule for making decisions under moral uncertainty depends crucially on the information provided by the moral theories in which one has credence. We showed how the resources from voting theory and social choice theory could be harnessed to help us develop such an account, arguing that in conditions of merely ordinal theories the Borda Rule was the correct account, in conditions of interval-scale measurability and intertheoretic incomparability variance voting was the correct theory, and in conditions of interval-scale measurability and intertheoretic comparability of choice-worthiness differences, *maximize expected choice-worthiness* (MEC) was the correct theory. We further argued that we can make sense of intertheoretic comparability, and that different moral theories are often comparable.

Finally, we showed how the Borda Rule, variance voting and MEC could be unified together in those situations (which will be the norm for real-life decision-makers) where the decision-maker faces different informational conditions all at once. We separate the theories into groups of theories that are mutually comparable with each other, and set the variance of choice-worthiness of each group to be equal (using Borda scores to represent the choice-worthiness of options on the ordinal theories) before taking an expectation. We suggested that this unified account could be thought of as an extension of *maximizing expected choice-worthiness*.

We then charted the implications of moral uncertainty for issues in metaethics and practical ethics. We argued that non-cognitivism has a very hard time providing a satisfactory account of moral uncertainty, and that the prospects of a positive solution look bleak. We argued that, though

moral uncertainty certainly has implications for practical ethics, those implications are not as obvious as has so far been presumed, and require at least somewhat controversial assumptions about what credences one ought to have in different moral theories. Also, we showed how the theory of decision-making under moral uncertainty gives us the resources to assess the value of gaining new moral information.

Though this book has covered a lot of ground, there is still much more to do. The topic of decision-making under moral uncertainty is as important, we believe, as the topic of decision-making under empirical uncertainty, and we believe that it should receive a commensurate amount of research attention. Right now, we have barely scratched the surface.

We will therefore suggest some promising and underexplored further lines of enquiry. There are some important topics that we simply didn't get to cover. These include the following.

- How to axiomatize decision-making under moral uncertainty.¹
- How we should assign deontic statuses, such as 'permissible' and 'impermissible', under moral uncertainty.
- What a reasonable credence distribution across different moral theories looks like.
- What the implications of moral uncertainty are for political philosophy, and in particular whether they can provide a justification for political liberalism.²

There are other topics that would certainly benefit from much greater study than we were able to give them. These include the following.

- How to make decisions under moral uncertainty given theories that posit incomparability between options.
- What grounds intertheoretic comparisons of value.
- What the most plausible intertheoretic comparison claims are between particular moral theories.
- The implications of moral uncertainty for practical ethics.

¹ In particular, building on the excellent work on evaluative uncertainty done by Riedener, 'Maximising Expected Value under Axiological Uncertainty'.

² For some work on this topic, see Evan Williams, 'Promoting Value as Such', *Philosophy and Phenomenological Research*, vol. 87, no. 2 (September 2013) pp. 392–416.

Finally, there are also promising lines of enquiry regarding other forms of normative uncertainty, in particular the following.

- Whether and how we ought to take decision-theoretic uncertainty into account in our decision-making.³
- Whether and how we ought to form beliefs in the face of epistemological uncertainty.⁴

We will end on a final, more speculative, note. In Chapter 9, we provided a framework for the value of gaining moral information. This framework can allow us to more clearly reflect on the value of moral philosophy as a whole.

What are the most important priorities that the world faces? When we ask that question, it's most natural to start comparing the magnitudes of some of the biggest known problems in the world—climate change, global poverty, disempowerment of women and minority groups—or to speculate on what might be major problems even if their status as such is still controversial, such as the suffering of wild animals, or the risk of human extinction. But it's plausible that the most important problem really lies on the meta-level: that the greatest priority for humanity, now, is to work out what matters most, in order to be able to truly know what are the most important problems we face.

The importance of doing this can hardly be overstated. Every generation in the past has committed tremendous moral wrongs on the basis of false moral views. Moral atrocities such as slavery, the subjection of women, the persecution of non-heterosexuals, and the Holocaust were, of course, driven in part by the self-interest of those who were in power. But they were also enabled and strengthened by the common-sense moral views of society at the time about what groups were worthy of moral concern. Given this dismal track record, it would be extremely surprising if we were the first generation in human history to have even broadly the correct moral worldview. It is of paramount importance, therefore, to figure out which actions society takes as common sensically permissible today we should really think of as

³ The first explorations of this idea are in MacAskill, 'Smokers, Psychos, and Decision-Theoretic Uncertainty'; Andrew Sepielli, 'What to Do When You Don't Know What to Do When You Don't Know What to Do...', *Noûs*, vol. 48, no. 3 (September 2014), pp. 521–44; and Michael G. Titelbaum, 'Rationality's Fixed Point (or: In Defense of Right Reason)', *Oxford Studies in Epistemology*, vol. 5 (2015), pp. 253–94.

⁴ To our knowledge there has been no sustained work on this topic, though the literature on peer disagreement and higher-order evidence sometimes veers into it.

barbaric. New moral information doesn't simply contribute some fixed amount of value to the world: insofar as it influences society's moral views, it has a multiplicative impact on all the value we might achieve into the future.

Given the importance of figuring out what morality requires of us, the amount of investment by society into this question is astonishingly small. The world currently has an annual purchasing-power-adjusted gross product of about \$127 trillion.⁵ Of that amount, a vanishingly small fraction—probably less than 0.05%⁶—goes to directly addressing the question: What ought we to do?

One might worry that, despite its importance and comparative neglect-ness, we simply cannot make meaningful progress on ethical questions, or that, if we do, our conclusions would have no influence anyway. But this, in our view, would be far too hasty. The impact of the median moral philosopher might be close to zero, but the mean is very high: on average, and over the long term, moral and political philosophy has made a huge difference to the world. Even just over the last few hundred years, Locke influenced the American Revolution and constitution, Mill influenced the woman's suffrage movement, Marx helped birth socialism and communism, and Singer helped spark the animal rights movement. If we broaden our horizons, and include Aristotle, Confucius, and Gautama Buddha in our comparison class, then it's hard to deny that the work of moral philosophy has shaped millennia of human history. And, simply by looking at the work in ethics done over the last few hundred years—by what is, globally speaking, a tiny number of people—it's hard not to believe that we have made significant progress.

⁵ International Monetary Fund, 'Report for Selected Country Groups and Subjects (PPP Valuation of Country GDP)', 2017, <https://goo.gl/RPV2Aw>.

⁶ 0.05% of annual gross world product is roughly \$60 billion, which we can regard as an upper bound estimate of the (very difficult to quantify) amount of investment that goes toward fundamentally normative questions. As a comparison, the total UK government expenditure in 2016 was £772 billion (David Gauke, *Budget 2016*, London: Stationery Office, 2016, p. 5), of which £300 million, or 0.04%, was spent on funding for the humanities and economic and social sciences through their research councils (Department for Business Innovation & Skills, *The Allocation of Science and Research Funding 2015/16*, May 2014, pp. 17, 23). Insofar as we should expect governments to spend more on normative research than the private sector, and rich countries to spend more than poor ones, and that the vast majority of humanities and social sciences funding goes to empirical research, it would be very surprising if the world as a whole invested a larger proportion of its resources into addressing normative questions than this.

We think, therefore, that considerations of moral uncertainty and the value of moral information should lead us to conclude that further normative research is one of the most important moral priorities of our time. Ideally, one day we will have resolved the deep moral questions that we face, and we will feel confident that we have found the moral truth. In the meantime, however, we need to do the best we can, given our uncertainty.

