Moral Ignorance and Bad Beliefs in Hostile Epistemic Environments

Agents frequently act from moral ignorance—they act on the basis of false beliefs about what morality requires they do. These agents, then, perform actions that are morally wrong, but would have been right had their relevant moral belief been true. Sometimes agents form false moral beliefs because they are epistemically irresponsible, failing to seek out evidence or deliberate properly given the evidence they do have. Other agents, however, form false moral beliefs because they exist in hostile epistemic environments. The paradigmatic example of such an epistemic environment is an *echo chamber*, a social epistemic structure wherein individuals share some set of beliefs and which serves to undermine disconfirming evidence and actively discredit those with opposing beliefs. Members of echo chambers who form false moral beliefs, then, do so largely because of their being systematically isolated from the broader epistemic environment.

There is substantial disagreement among philosophers over whether an agent's being morally ignorant—holding a false moral belief—exculpates or otherwise mitigates their blameworthiness for acting wrongly from that ignorance. A view on which moral ignorance does not exculpate holds that an agent who holds false moral beliefs as a result of existing in a bad epistemic environment, like an echo chamber, is blameworthy for the wrong actions they perform on the basis of those beliefs. Now, this might seem like the right result, for some of the most abhorrent and harmful beliefs arise in and are reinforced by echo chambers—think conspiracy theories, neo-Nazism, religious extremism. My aim, however, is to argue that this verdict is delivered prematurely and without recognizing the full moral and epistemic import of echo chambers.

In this talk, I explore two ways that consideration of hostile epistemic environments may challenge this verdict. The first route considers whether beliefs formed in echo chambers and other bad epistemic environments can be epistemically justified. For perhaps it's unreasonable to blame someone for acting on the basis of a belief when they (epistemically) ought not have believed otherwise. This might be the case if, as some philosophers have suggested, it's rational to remain in one's echo chamber. Whether an agent in an echo chamber is blameworthy for performing a wrong action, then, will turn on whether beliefs formed in echo chambers are epistemically justified. Asking this question, however, overlooks a more fundamental one, bringing us to the second route. We should consider what it's possible for these agents to believe: can agents in echo chambers be *reasonably expected* to form beliefs that diverge from those prescribed by their epistemic communities? Drawing on Ngyuen's (2020) idea that escaping an echo chamber requires a complete epistemic reboot, requiring an agent to discard all of their beliefs and start over, I suggest that we should give a negative answer. Once an agent is in an echo chamber, it is unreasonable to think that they could have believed anything but what they did. We might think that this means we should be hesitant to ascribe blameworthiness to these agents, and perhaps we should be. The more important upshot, however, is that moral evaluations of our beliefs and the actions we perform on their basis cannot be determined independent of the epistemic environments in which we exist and that we cannot escape.

Apologies, blame, and taking responsibility for ignorance

People often say, 'don't apologise, you didn't know any better'. For example, I might be in a country where it's polite to take one's shoes off before entering another person's home. If I walk through someone's home with my shoes on, it would be right for them to tell me off and it would be natural for me to take my shoes off and start to apologise for my action. However, upon finding out that I was ignorant of the shoe etiquette, it would also be natural for my hosts to say, 'don't apologise, you didn't know better'. This suggests that ignorance can make it the case that it's not fitting to apologise; because I did not know that it was impolite to wear shoes inside someone's home, it's not fitting for me to apologise for doing so. However, it's not clear that all cases of ignorance are cases in which it would not bet fitting to apologise. Imagine that you are allergic to penicillin, but your doctor administers penicillin because they didn't know about your allergy. Despite your doctor's ignorance, it would still be fitting for your doctor to apologise for their mistake. This leaves us with a question: when is it fitting to apologise for actions performed under ignorance, and when it is not? The answer is that it's never fitting to apologies for *actions* perform under ignorance, but that it can be fitting to apologise for the *ignorance* itself. As members of a community, we take on roles that in result in our having a responsibility know certain things. In cases where you have a responsibility to know what you were ignorant of, a failure to do so makes it fitting to apologies for your ignorance. Importantly, in cases of ignorance it will only be fitting to apologise for your ignorance, and it will not be fitting to apologise for the resulting action(s).

The paper will go as follows: In Section 1, I will give an overview of what apologies are. Section 2 will look at cases of non-normative ignorance. Here we will see that sometimes intuitions point to apologies being fitting in cases of ignorance, while other times it seems right to say that it is not fitting to apologise. A similar clash of intuitions can be found in cases where a *prima facie* duty has been violated in order for someone to perform an all things considered right action, and in cases where someone has performed a wrong action but done so blamelessly. Given the similarities between cases, Section 3 turns to explore the fittingness of apologies in cases of violated *prima facie* duties and blameless wrongs. The findings of Section 3 will show that in cases of ignorance it's fitting to apologise when you had a responsibility to know what you were ignorant of. Section 4 will explore this further by unpacking when we do have a responsibility to know. Section 4 will end by arguing that in cases of ignorance, while it is fitting to apologise for being ignorant, it is not fitting to apologise for the resulting action. Section 5 will conclude the paper by looking at cases of normative ignorance.

Moral uncertainty, objectivism, and epistemic access

Almost all of us are forced to make ethical and political decisions despite being at least somewhat uncertain about which moral theory is correct. Thus, almost all of us stand in need of some version of decision theory which can recommend courses of action to morally uncertain decision makers. Existing proposals in the literature include:

- 'Maximizing Expected Choiceworthiness' (MEC);
- following 'My Favourite Theory' (MFT); or 'Maximising Expected Truthlikeness' (MET);
- several voting- and bargaining-theoretic approaches, inspired by social choice theory.

All of these proposals are standardly defined in terms of '*choiceworthiness*' – where the choiceworthiness of some option A according to any given moral theory T is defined as the strength of the decision maker's all-things-considered moral reason in favour of choosing A according to T.

Sadly, *objectivist* moral theories pose an immediate (and seriously neglected) problem for all of these choiceworthiness-based proposals. According to objectivism, the choiceworthiness of an option depends upon the *actual* features and consequences of that option, regardless of whatever the decision maker happens to *believe* about those features and consequences. For instance, objective consequentialism says that the choiceworthiness of any given option depends upon the consequences that would *actually* result from choosing that option, even if our decision maker is highly descriptively uncertain about the content of those consequences.

The problem for all of the choiceworthiness-based responses to moral uncertainty is that a decision maker with nontrivial credence in any objectivist moral theory may be unable to determine which option is recommended by these proposals. In other words, these proposals will often fail to be properly 'action-guiding.'

My aim in this paper is to consider whether and how each of these proposals can be amended to address this problem of objectivism. Firstly, I will argue that an elegant modification to MEC is just to propose that a morally uncertain decision maker should calculate the expected choiceworthiness of each option using her probability distribution over '*complete* theories' – where a complete theory is a theory of all of the moral *and* descriptive facts. It has been suggested (e.g. by Tarsney) that this proposal won't respect the risk attitudes of certain non-objectivist moral theories. However, I demonstrate that this objection is applicable only to approaches that calculate expected choiceworthiness over all of the possible *beliefs* that our decision maker could have in a complete theory, as opposed to over all of the possible complete *states of affairs* which could obtain.

Unfortunately, this complete-theories modification strategy will not work for MFT, MET, and the voting- or bargaining-theoretic proposals, since this strategy will have strange implications for agents with high credence in an objective moral theory, but who are nonetheless highly uncertain descriptively.

Could we instead redefine these non-MEC proposals in terms of the '*decision*worthiness' of options under different possible *decision procedures*? I will argue that this approach is not particularly promising either. The overall upshot of the paper is that rival proposals to MEC face an unresolved problem in handling objectivist moral theories.