

ARTICLE

Practical Ethics Given Moral Uncertainty

William MacAskill

University of Oxford

Corresponding author. Email: william.macaskill@philosophy.ox.ac.uk

(Received 15 September 2017; revised 16 January 2019; accepted 19 January 2019)

Abstract

A number of philosophers have claimed that we should take not just empirical uncertainty but also fundamental moral uncertainty into account in our decision-making, and that, despite widespread moral disagreement, doing so would allow us to draw robust lessons for some issues in practical ethics. In this article, I argue that, so far, the implications for practical ethics have been drawn too simplistically. First, the implications of moral uncertainty for normative ethics are far more wide-ranging than has been noted so far. Second, one can't straightforwardly argue from moral uncertainty to particular conclusions in practical ethics, both because of 'interaction' effects between moral issues, and because of the variety of different possible intertheoretic comparisons that one can reasonably endorse.

I. Introduction

Recently, a number of philosophers have argued that (i) we should take not just empirical uncertainty but also fundamental moral uncertainty into account in our decision-making, and also that (ii) doing so has significant implications for practical ethics.¹

This literature has principally focused on the topics of abortion and vegetarianism. The argument runs approximately as follows. Consider, first, the following case of decision-making under empirical uncertainty:

Speeding

Julia is considering whether to speed round a blind corner. She thinks it's pretty unlikely that there's anyone crossing the road immediately around the corner, but she's not sure. If she speeds and hits someone, she will certainly severely injure them. If she goes slowly, she certainly will not injure anyone, but will get to work slightly later than she would have done had she sped.

¹Alexander A. Guerrero, 'Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution', *Philosophical Studies* 136 (2007), pp. 59–97; T. Lockhart, *Moral Uncertainty and its Consequences* (New York, 2000); Graham Oddie, 'Moral Uncertainty and Human Embryo Experimentation', *Medicine and Moral Reasoning*, ed. K. W. M. Fulford, G. Gillett and J. Martin Soskice (Cambridge, 1994), pp. 144–61; Dan Moller, 'Abortion and Moral Risk', *Philosophy* 86 (2011), pp. 425–43.

Her decision-situation looks as follows:

	Someone crossing	No-one crossing
Speed	Significant wrong	Permissible
Go Slow	Permissible	Mild personal cost

In this situation, both expected value reasoning and common sense recommend that Julia should not speed.

But if we agree with this in a case of purely empirical uncertainty, and we think that we should treat moral and empirical uncertainty analogously, then we should in general think that it’s impermissible to eat meat.² Consider the following case:

Vegetarianism

Harry is considering whether to eat meat or a vegetarian option for dinner. He thinks it’s pretty unlikely animals matter morally, but he’s not sure. If he eats meat and animals do matter morally, then he commits a grave wrong. If he eats the vegetarian option, he will certainly not commit a grave wrong, though he will enjoy the meal less than he would have done had he eaten meat.

	Animals matter	Animals don’t matter
Eat meat	Significant wrong	Permissible
Eat vegetarian	Permissible	Mild personal cost

Here, the decision situation is analogous to the decision situation in *Speeding*. Even if Harry is highly confident in the view that animals don’t matter, his credence in the view that they do matter generates a significant risk of doing something gravely wrong, outweighing the greater likelihood of missing out on a mild prudential benefit. If we thought that Julia shouldn’t speed in *Speeding*, then we should think that Harry shouldn’t eat meat in *Vegetarianism*.

A similar argument can be made for abortion. Consider the following case:

Abortion

Isobel is twenty weeks pregnant and is considering whether to have an abortion. She thinks it’s pretty unlikely that twenty-week-old fetuses have a right to life, but she’s not sure. If she has an abortion and twenty-week-old fetuses do have a right to life, then she commits a grave wrong. If she has the child and gives it up for adoption, she will certainly not commit a grave wrong, though she will bear considerable costs as a result of pregnancy, childbirth and separation from her child.

²See Guerrero, ‘Don’t Know, Don’t Kill’ and Moller, ‘Abortion and Moral Risk’. Sometimes this and the case against abortion are presented as a *dominance* argument, where vegetarianism, or having a child, is suggested to be certainly permissible (Lockhart, *Moral Uncertainty*, ch. 2; Brian Weatherston, ‘Review of Ted Lockhart, *Moral Uncertainty and Its Consequences*’, *Mind* 111 (2002), pp. 693–6). However, I think that we should be considering how to make decisions in light of all the possible reasons for action that one has. And if one believes that there is no moral reason against eating meat, whereas there is a prudential reason in favour of eating meat, then eating meat is the most all-things-considered choice-worthy option. So the ‘dominance’ form of the argument will almost never apply.

	Foetuses have a right to life	Foetuses have no right to life
Have abortion	Very gravely wrong	Permissible
Give up for adoption	Permissible	Significant personal cost

In this case, the prudential cost to the decision-maker is higher than it is in *Speeding* or *Vegetarianism*. But the potential moral wrong, if the view that foetuses have a right to life is correct, is also much greater. So, again, it seems that even if Isobel is fairly confident in the view that foetuses have no right to life, the risk that they do is sufficient to outweigh the significant prudential reason in favour of having the abortion, and Isobel should not have the abortion.

If this argument works, then it is like the philosopher's stone for practical ethicists: it would mean that we could draw robust lessons for practical ethics even despite extensive disagreement among moral philosophers. As Ted Lockhart comments: 'The significance of this argument is that, if sound, it shows that much of philosophers' discussion of the morality of abortion is for practical (i.e., decision-making) purposes unnecessary.'³ Some philosophers endorse the implications of moral uncertainty for vegetarianism and abortion;⁴ others take them to be a *modus tollens*.⁵ But all authors so far seem to agree that taking moral uncertainty into account in one's decisions really does have these implications, does so in a rather straightforward way, and does so largely independently of the credences that one has in different moral views, as long as those credences are broadly reasonable.

In this article, I'm going to make things more complicated, in two ways. First, I show that the *prima facie* implications of moral uncertainty for issues in normative ethics are far more wide-ranging than has been noted in the literature so far.

Second, I show how one can't straightforwardly argue from moral uncertainty to particular conclusions in practical ethics, using abortion and vegetarianism as particular examples. I argue for this on two grounds: first, because of 'interaction' effects between moral issues; and, second, because of the variety of different possible intertheoretic comparisons that one can reasonably endorse. The conclusion I reach is that, before drawing out conclusions from moral uncertainty based arguments, one first has to do the difficult job of figuring out what one's credences in different moral viewpoints are or ought to be.

To be clear, I don't take the implications I draw to be problems for the 'maximize expected choice-worthiness' (MEC) account of decision-making under moral uncertainty (where the 'choice-worthiness' of an option represents the all-things-considered

³Lockhart, *Moral Uncertainty*, p. 52.

⁴Moller: '[the moral uncertainty argument] does seem to suggest, however, that there is a moral reason – probably not a weak one – for most agents to avoid abortion' ('Abortion and Moral Risk', p. 443). Lockhart: 'In the vast majority of situations in which decision-makers decide whether to have abortions, *not* having an abortion is the *reasonable* choice of action' (*Moral Uncertainty and its Consequences*, p. 52). Pope John Paul II: 'the mere probability that a human person is involved [in the practice of abortion] would suffice to justify an absolute clear prohibition of any intervention aimed at killing a human embryo' (http://w2.vatican.va/content/john-paul-ii/en/encyclicals/documents/hf_jp-ii_enc_25031995_evangelium-vitae.html).

⁵Weatherson: '[implications] so striking we might fear for its refutation by a quick *modus tollens*' ('Review of Lockhart', p. 694). Guerrero: '[maximizing expected moral value] is not the reading that I prefer, in part because of cases like [abortion]' ('Don't Know, Don't Kill', p. 91).

strength of reasons in favour of that option).⁶ Rather, all I aim to show is that, so far, MEC's implications for topics in normative ethics have been understood in too simple-minded a way. Indeed, for the purposes of this article, I will assume that MEC is approximately correct. That is, the appropriate option under moral uncertainty is the one with the highest expected choice-worthiness, where the expected choice-worthiness of an option is given by the sum, over all moral theories, of the decision-maker's credence in that theory multiplied by the choice-worthiness of the option, if that theory is true. (I will use the technical term 'appropriate' to refer to what is the correct choice under moral uncertainty, in order to distinguish this concept from what one ought to do according to the true moral theory.) Note that one can distinguish two forms of MEC depending on whether the credences that are referred to are the decision-maker's *actual* degrees of belief, or the degrees of belief that the decision-maker *ought* to have. Though this is an important issue, it is one I will leave to the side in this article; I'll assume throughout that the decision-maker has reasonable credences given her epistemic state.

There are known problems for the maximizing expected choice-worthiness account,⁷ which I will have to put to the side. I will assume that all moral views under consideration provide an interval-scale measurable notion of choice-worthiness (such that it's meaningful to speak of ratios of differences in choice-worthiness between options), and that quantities of choice-worthiness can be meaningfully compared across different moral views (that is, I assume that there are meaningful intertheoretic comparisons of choice-worthiness). I will also assume that we are not considering moral views to which we assign very low credence – we might rightfully balk at the idea of maximizing expected choice-worthiness when we are dealing with very low probability theories that posit astronomically high stakes. Finally, I assume that all moral theories we consider are complete: there is no incomparability of value. I acknowledge that this is a significant limitation, but there are major open questions regarding what it is rational to do when facing incomparability under moral uncertainty,⁸ and a theory of how to handle this is a task for an article all by itself.

In a full analysis of the practical implications of moral uncertainty, all these factors would be taken into account.⁹ However, philosophers have yet to understand the practical implications of moral uncertainty even with these simplifying assumptions; the task of understanding moral uncertainty's implications for practical ethics without these simplifying assumptions will therefore have to be a task for further work.

II. Implications for normative ethics

Though the moral uncertainty literature so far has focused on vegetarianism and abortion, there are many issues in normative ethics where there appear to be clear *prima facie* implications of taking moral uncertainty into account in our decision-making, most of which have not yet been noticed.¹⁰ This section provides a brief overview of them.

⁶For this account, see Andrew Sepielli, 'What to Do When You Don't Know What To Do', *Oxford Studies in Metaethics* 4 (2009), pp. 5–28; William MacAskill, 'Moral Uncertainty' (DPhil dissertation, Oxford University, 2014).

⁷Jacob Ross, 'Rejecting Ethical Deflationism', *Ethics* 116 (2006), sects. 4–5.

⁸See William MacAskill, 'The Infectiousness of Nihilism', *Ethics* 123 (2013), pp. 508–20.

⁹In 'Normative Uncertainty as a Voting Problem', *Mind* 125 (2016), pp. 967–1004, I suggest an extension of maximizing expected choice-worthiness that attempts to deal with some of these problems.

¹⁰The implications of moral uncertainty have been discussed for abortion (Lockhart, *Moral Uncertainty*, ch. 3; Moller, 'Abortion and Moral Risk'), embryo destruction (Oddie, 'Moral Uncertainty and Human

Beneficence

Consider Peter Singer's argument that failing to donate to the developing world is as wrong, morally, as letting a child drown in front of you.¹¹ If one has even a moderate credence in that view, then it seems that under moral uncertainty it's appropriate to donate a substantial proportion of one's resources to save the lives of strangers. Not-donating involves a risk of doing something as wrong as letting a child drown in front of you; whereas donating involves only the risk of needlessly incurring a moderate prudential cost. The situation therefore seems analogous to *Speeding*: for someone who is unsure about whether Singer's arguments work, it would be inappropriate not to donate.

A distinct argument for the same conclusion can be gained by considering the acts/omissions distinction. Even if you are fairly confident in the moral relevance of the distinction between acts and omissions, you shouldn't be completely certain in that view. You should give some credence to the idea that it's just as wrong to let someone die as it is to intentionally kill them. In which case, you should have some credence that letting distant strangers die because of failing to donate to effective non-profits is roughly as wrong as actively killing them. This gives a second argument for why considerations of moral uncertainty provide an argument for donating a substantial proportion of your resources to save the lives of strangers.

Partiality

Under moral uncertainty, one should give some extra moral weight to one's friends and family's interests, even if your preferred moral view is impartial.

Even if you are confident that the well-being of your friends and family is equally as valuable as the well-being of distant strangers, you should not be certain in that view: you should have some credence that the well-being of your friends and family is more valuable than the well-being of distant strangers. However, you should have almost no credence that the well-being of distant strangers is *more* important than the well-being of your friends and family. So you should therefore give the interests of your friends and family some extra weight, though not as much weight as if you were completely convinced of the partialist moral view. If you could benefit your friend or a stranger by the same amount, it's therefore appropriate to benefit your friend over the stranger.

Prioritarianism, equality, utilitarianism

Under moral uncertainty, you should treat benefits to the badly off as being more important than providing the same benefits to the well off, even if you are fairly confident that they should be treated in the same way. The argument for this is analogous to the argument I just made about partiality. You should have some credence in the view that it's more important to give a benefit of a given size to someone who is worse off

Embryo Experimentation'), vegetarianism (Moller, 'Abortion and Moral Risk', pp. 426, 441–3; Guerrero, 'Don't Know, Don't Kill', pp. 76–82), the ethics of killing more generally (Guerrero, 'Don't Know, Don't Kill') and duties of beneficence (Lockhart, *Moral Uncertainty* ch. 5; Weatherston, 'Review of Lockhart'). I don't know of other examples of the practical issues being discussed, so believe that the suggested implications for partiality, egalitarianism, the suffering/happiness trade-off, theories of well-being, welfarism, egoism and population ethics are novel.

¹¹Peter Singer, 'Famine, Affluence and Morality', *Philosophy & Public Affairs* 1 (1972), pp. 229–43.

rather than to someone who is better off; this view is entailed by both prioritarianism and egalitarianism. In contrast, you should have almost no credence in the view that one ought to give a benefit of a given size to someone who is better off rather than worse off: this is not entailed by any reasonable moral position. So, under moral uncertainty, it will be appropriate to give a benefit of a given size to someone who is worse off rather than someone who is better off.

Alleviation of suffering

Under moral uncertainty you should treat alleviating suffering as more important than increasing happiness. Again, the reasoning is analogous to our last two arguments. According to some plausible moral views, the alleviation of suffering is more important, morally, than the promotion of happiness. According to other plausible moral views (such as classical utilitarianism), the alleviation of suffering is equally as important, morally, as the promotion of happiness. But there is no reasonable moral view on which the alleviation of suffering is less important than the promotion of happiness. So, under moral uncertainty, it's appropriate to prefer to alleviate suffering rather than to promote happiness more often than the utilitarian would (though less often than would the view that preventing suffering is more important than alleviating happiness).

Theories of well-being

Some theories of well-being claim that having 'objective' goods, like knowledge or appreciation of beauty, intrinsically make a person's life go better; whereas other theories, such as hedonism and preference-satisfactionism, do not place value on those goods beyond how they contribute to positive mental states or to preference-satisfaction. But no theories of well-being claim that possessing objective goods intrinsically makes a person's life go worse. (Of course, some theories of well-being may imply, together with the relevant empirical facts, that possessing certain objective goods sometimes *instrumentally* makes a person's life worse. But I'm here making a claim about intrinsic rather than instrumental goodness.)

Under uncertainty about theories of well-being, therefore, one should treat alleged objective goods as having some positive intrinsic value, but not as much as they have on the objective list theory of well-being.

Welfarism

Similarly, some views, such as utilitarianism, place value only on people's welfare. On other views, there are non-welfarist goods that have intrinsic value, such as great works of art or a well-preserved natural environment. But on no reasonable moral view are the supposed non-welfarist goods of negative intrinsic value. So, if you are unsure between welfarism and non-welfarist views, then under moral uncertainty you should treat the alleged non-welfarist goods as having some intrinsic value, but not as much value as the non-welfarist view regards them as having.

Egoism and altruism

Given egoism, you only have reasons to improve your own welfare. On other moral views, you also have intrinsic reasons to improve the lives of others or respect their

rights. But on no plausible moral views is it the case that you have intrinsic reasons to harm others, or violate their rights. So, if you are uncertain between egoism and other moral views, then one has some reasons to benefit others, though not quite as strong reasons as you would have if you did not believe egoism at all.

Population ethics

Extending moral uncertainty to issues of population ethics has three main implications, concerning total versus critical-level views, separable versus non-separable views, and person-affecting versus non-person-affecting views.¹²

First, let us consider only separable non-person-affecting views: that is, views on which the value of adding an additional person to the population is independent of how many other people already exist, who they are, and what their well-being levels are. Among such views, there are two plausible theories: the *total* view, according to which the goodness of bringing a new person into existence is given by how much better or worse that person's life is than a 'neutral life', and *critical-level* views, according to which it's good to bring into existence a person if their life is above a certain level of well-being c , neutral if their life is at level c , and bad if their life is below c .¹³

Under uncertainty between the total view and critical-level views, bringing a new person into existence would have positive expected choice-worthiness if their lifetime welfare is above an 'expected' critical-level c^* , where c^* is lower than the critical-level views one has credence in would claim, but greater than 0. This is because no plausible critical-level view endorses a negative critical-level, since such a view would imply that bringing into existence lives with negative welfare has positive value. Given that the total view is equivalent to a critical-level view with the critical-level set to zero, the critical-levels over which we are uncertain go from 0 to a positive number, and the 'expected' critical-level must fall within this range.¹⁴

Second, let us consider uncertainty over separable and non-separable views. Non-separable views include the average view, according to which the goodness of a population is given by the average well-being of that population, and views according to which the goodness of a population is determined by both the average well-being of the population and the total well-being of the population.¹⁵ Under uncertainty between separable views and non-separable views, one will place weight on both the average well-being of the population (or other 'quality' measures) and the sum total of well-being that is above c^* minus the total well-being that is below c^* .

Finally, we turn to uncertainty between person-affecting and non-person-affecting views. On non-person-affecting views, bringing a new person with a life worth living into existence is either good or bad, depending on the well-being level of the person in question. But on

¹²For a more comprehensive discussion of these different views, see Hilary Greaves and Toby Ord, 'Moral Uncertainty about Population Axiology', *Journal of Ethics and Social Philosophy* 12 (2017), pp. 135–67.

¹³See Charles Blackorby and David Donaldson, 'Social Criteria for Evaluating Population Change', *Journal of Public Economics* 25 (1984), pp. 13–33; Charles Blackorby, Walter Bossert and David Donaldson, 'Intertemporal Population Ethics: Critical-Level Utilitarian Principles', *Econometrica* 63 (1995), pp. 1303–20; J. Broome, *Weighing Lives* (Oxford, 2004); C. Blackorby, W. Bossert and D. Donaldson, *Population Issues in Social Choice Theory, Welfare Economics, and Ethics* (Cambridge, 2005).

¹⁴This idea is developed in Greaves and Ord, 'Moral Uncertainty about Population Axiology'.

¹⁵See Thomas Hurka, 'Value and Population Size', *Ethics* 93 (1982), pp. 496–507; Yew-Kwang Ng, 'What Should We Do About Future Generations?', *Economics and Philosophy* 5 (1989), pp. 235–53.

person-affecting views, there is no reason for or against bringing someone with a life worth living into existence.¹⁶ So a decision-maker with uncertainty over person-affecting and non-person-affecting views will, in general, think that it's either good or bad to bring a particular person into existence; the decision-maker's credence in person-affecting views has no effect on that issue. Making this more precise, let us use n to refer to the level of well-being at which, conditional on non-person-affecting views, one would think it neither good nor bad, in terms of expected choice-worthiness, to bring a new person with that well-being level into existence. Because person-affecting views provide no reason either way, when it comes to lives that are worth living, n is the level above which it is of positive expected choice-worthiness to bring someone into existence (and below which it is of negative expected choice-worthiness) no matter one's credence in person-affecting views, as long as one has some credence in non-person-affecting views.

III. Interaction effects

As I noted at the outset, some philosophers have suggested that the implications of maximizing expected choice-worthiness are so clear on some issues in practical ethics that we can cease further work on the normative question of which view on the issue is the correct one.¹⁷

I believe that to be a mistake. So far, commentators haven't noticed just how broad the range of different implications of moral uncertainty based arguments is. That is obviously an oversight in so far as it means they've underestimated the importance of moral uncertainty based reasoning. But it's also an oversight in so far as it impacts on how moral uncertainty based arguments should be applied, including in the central examples of vegetarianism and abortion. We cannot simply look at how moral uncertainty impacts on one debate in practical ethics in isolation; moral uncertainty arguments have very many implications for practical ethics, and many of those interact with one another in subtle ways.

Consider vegetarianism. Moller states that 'avoiding meat doesn't seem to be forbidden by any view. Vegetarianism thus seems to present a genuine asymmetry in moral risk: all of the risks fall on the one side.'¹⁸ Similarly, Weatherson comments that, 'the actions that Singer recommends ... are certainly morally permissible ... One rarely feels a twang of moral doubt when eating tofu curry.'¹⁹

That is, the moral uncertainty argument for vegetarianism got its grip because there was supposedly no or almost no moral reason in favour of eating meat. Once we consider *all* the implications of moral uncertainty, however, this is no longer true.

We saw that, given moral uncertainty, it's good (in expectation) to bring into existence beings with lives that are sufficiently good (above the expected critical-level c^*). And some types of animals raised for consumption, such as cows, sheep, humanely raised chicken and pork, plausibly have lives that are worth living.²⁰ Depending on

¹⁶See Jan Narveson, 'Moral Problems of Population', *The Monist* 57 (1973), pp. 62–86.

¹⁷For example, Lockhart, *Moral Uncertainty*, p. 52.

¹⁸Moller, 'Abortion and Moral Risk', p. 441.

¹⁹Weatherson, 'Review of Lockhart', p. 693.

²⁰An assessment of the welfare levels of various farm animals is given in F. B. Norwood and J. L. Lusk, *Compassion, by the Pound: The Economics of Farm Animal Welfare* (New York, 2011), p. 229, and in N. Cooney, *Veganomics: The Surprising Science on What Motivates Vegetarians, from the Breakfast Table to the Bedroom* (New York, 2013), ch. 1.

exactly how one distributes one's credences across total views and critical-level views, one might reasonably judge that these lives are above the expected critical-level c^* .

Importantly, when you choose to buy meat you aren't killing animals. Instead you are increasing demand for meat, which incentivizes farmers to raise (and then kill) additional animals. By buying and eating cows, sheep, free range chicken and pork, you cause animals with possibly fairly happy lives to come into existence which would not otherwise have lived. On some mainstream consequentialist views (such as total utilitarianism), it's therefore wrong *not* to purchase the meat of such animals.

Our decision situation is therefore more complicated than commentators have suggested. Let's call the view that regards eating meat if animals matter as significantly wrong (perhaps because of complicity in an immoral act²¹) the *non-consequentialist view*, and the view that eating meat is right or wrong only in so far as it impacts upon creatures' welfare as the *consequentialist view*. The relevant decision situation could therefore be represented as follows:

	Animals matter		Animals don't matter
	non-consequentialist view	consequentialist view	
Eat meat	Significant wrong	Permissible	Permissible
Eat vegetarian	Permissible	Significant wrong	Mild personal cost

Importantly, this means that we can't state that, given moral uncertainty, one ought, given any reasonable set of moral credences, to be vegetarian. It might be that you find the total view of population ethics very plausible, in which case eating beef and lamb might have higher expected choice-worthiness than eating vegetarian. Alternatively, you might find the total view of population ethics very implausible, but find the idea that you shouldn't be complicit in immoral actions very plausible; in which case, under moral uncertainty vegetarianism might indeed be the more appropriate course of action. It all depends on controversial conclusions about how confident you should be in different normative theories.

One might respond by restricting the scope of the argument. Rather than claiming that moral uncertainty considerations lead to vegetarianism, one might instead argue that they entail simply not eating those animals (for example, factory farmed chickens) whose lives have been so bad as not to be worth living. In this case, the argument that eating meat is good because it brings into existence animals with happy lives would not go through; eating this meat brings into existence animals with net unhappy lives which, almost everyone would agree, is a bad thing to do. This, one might argue, is still an example where, as Lockhart suggests, philosophers' discussion is unnecessary for practical purposes.

But, even here, I do not think that one can draw conclusions from moral uncertainty based arguments without also invoking at least somewhat controversial assumptions about what credences one ought to have in different moral views. First, it's a question

²¹For example, this view can take the form of an endorsement of Christopher Kutz's Complicity Principle (*Complicity: Ethics and Law for a Collective Age* (Cambridge, 2000), p. 122), according to which 'I am accountable for the harm or wrong we do together, independently of the actual difference I make', combined with the claim that we together wrong the animals we eat.

for moral philosophy (in part) what animals have lives that are and aren't worth living; it's not a wholly unreasonable view that even factory farmed chickens have lives that are worth living. If that were true, then there would be at least one moral view according to which one ought to eat factory farmed chicken. In order to make moral uncertainty based arguments entail not-eating factory farmed chicken, one must argue (at least somewhat controversially) that those moral views according to which factory farmed chickens do not have lives worth living are significantly more plausible than those moral views according to which they have lives that are worth living.

Moreover, remember that consideration of moral uncertainty seemed to show that we have strong duties of beneficence to help the global poor. Restricting your diet costs time and money, which could be used fighting poverty, saving lives in the developing world. Over the course of your life, you could probably save enough time and money by allowing yourself to (sometimes) eat meat to save a life in the developing world.²² On some views concerning the relative value of human and non-human lives, this implies that it is permissible to eat these animals, and moderately wrong not to eat them. On other views, the implications would be reversed. This means that a more accurate representation of the decision situation looks as follows:

	Animals matter a lot		Animals matter a little		Animals don't matter	
	Obligation to donate	No obligation to donate	Obligation to donate	No obligation to donate	Obligation to donate	No obligation to donate
Eat meat	Significant wrong	Significant wrong	Significant wrong	Moderate wrong	Significant wrong	Permissible
Eat vegetarian	Permissible	Permissible	Moderate wrong	Permissible	Significant wrong	Mild personal cost
Eat cheapest & donate	Moderate wrong	Moderate wrong	Permissible	Moderate wrong	Permissible	Mild personal cost

Again, therefore, we can no longer argue that maximizing expected choice-worthiness would recommend eating vegetarian no matter what reasonable credences one has across moral views. Rather, what conclusion we reach depends on substantive views about (i) how plausible different moral views are; (ii) the strengths of your obligations, if those views are correct.

Similar considerations apply to abortion in so far as one believes that, at the time of abortion, the embryo or foetus is not yet a person. First, even though on ordinary

²²According to the latest estimates from GiveWell, it costs about \$3,200 to do the equivalent amount of good to saving a life in poor countries ('GiveWell cost-effectiveness analysis', November 2016, https://docs.google.com/spreadsheets/d/1KiWfiAGX_QZhRbC9xkzf3I8IqsXC5kkr-nwY_feVlcM). In order for the costs of a strict vegetarian diet to be greater than the cost to save a life, the strict vegetarian diet would only have to cost an additional \$1.53 per week over a span of 40 years. One might object that a vegetarian diet is cheaper than an omnivorous diet. This may, typically, be true. However, because one loses options by being vegetarian, a vegetarian diet must be at least as costly as the diet one has if one acts on the maxim 'eat whatever's cheapest', and it seems unlikely that such a maxim would *never* involve eating meat.

morality, the decision whether to have a child is of neutral value, on some other theories this is not the case. In particular, on some moral views, it is wrong to bring into existence even a relatively happy child. On person-affecting views there is no reason in virtue of the welfare of the child to have a child; and if you believe that the world is currently overpopulated, then you would also believe that there are moral reasons against having an additional child. On critical-level views of population ethics, it's bad to bring into existence lives that aren't sufficiently happy; if the critical-level is high enough, such that you thought that your future child would probably be below that level, then according to a critical-level consequentialist view you ought not to have the child. On environmentalist or strong animal welfare views it might be immoral to have a child, because of the environmental and animal welfare impact that additional people typically have. Finally, on anti-natalist views, the bads in life outweigh the goods, and it's almost always wrong to have a child.

This means, again, that we cannot present the decision of whether to have an abortion given moral uncertainty as a decision where one option involves some significant moral risk and the other involves almost no moral risk. We should have at least some credence in all the views listed in the previous paragraph; given this, in order to know what follows from consideration of moral uncertainty we need to undertake the tricky work of determining what credences we should have in those views. (Of course, we would also need to consider those views according to which it's a good thing to bring into existence a new person with a happy life, which might create an additional reason against having an abortion.)

Moreover, as with the case of vegetarianism, we must consider the issue of opportunity cost. Carrying a child to term and giving it up for adoption costs time and money (in addition, potentially, to psychological distress) that could be used to improve the lives of others. According to a pro-choice view that endorses Singerian duties of beneficence, one would be required to have an abortion in order to spend more time or money on improving the lives of others. Again, what seems appropriate under moral uncertainty is critically dependent on what exactly the decision-maker's credences across different moral theories are.

In the above examples, we have just looked at the interaction effects between vegetarianism and abortion and duties of beneficence and population ethics. But, as noted in the previous section, there are *very many* implications of taking moral uncertainty into account. The interactions between these various implications may be quite subtle; a full analysis of the implications of moral uncertainty for any particular topic in practical ethics would need to take *all* of these implications into account.

IV. Intertheoretic comparisons

Interaction effects are one way in which the alleged implications of moral uncertainty might not follow. Choice of intertheoretic comparisons is another.

Consider vegetarianism again. Let's (simplistically) suppose that on the ordinary morality view, the welfare of non-human animals has one ten thousandth the moral weight of the welfare of humans, and that on the 'all animals are equal' view, the welfares of humans and animals are of equal moral worth. (Of course, few moral views regard the welfare of animals and humans as equal. This does not matter for the purposes of the example.) When philosophers have argued from moral uncertainty to vegetarianism, they've implicitly invoked one specific way of making intertheoretic comparisons between the 'ordinary morality' view and the 'animal welfare' view. But

that isn't the only way of making the comparison. Here are two different ways of making the intertheoretic comparison:²³

Option	Ordinary Morality	All-Animals-Are-Equal-1	All-Animals-Are-Equal-2
1 unit of human welfare	10,000	10,000	1
1 unit of animal welfare	1	10,000	1
0 units of welfare	0	0	0

There are in fact two natural ways of revising the ordinary morality view in order to make the welfare of all animals equal. On the first view, All-Animals-Are-Equal-1, the revision is that animal welfare is much more valuable than the ordinary morality view supposes. On the second view, All-Animals-Are-Equal-2, the revision is that human welfare is much less valuable than the ordinary morality view supposes.

I believe that both ways of making the intertheoretic comparison are 'permissible': they represent different theories, one may have credence in either, and the question of what credence one ought to have in the different comparisons is largely a question for normative ethical theorizing. But whether or not the moral uncertainty based argument for vegetarianism goes through depends to a large extent on which of these two intertheoretic comparisons we invoke. If Harry (in the original example) is unsure between Ordinary Morality and All-Animals-Are-Equal-1, then it is indeed true that he risks a grave wrong by eating meat. If, however, he is unsure between Ordinary Morality and All-Animals-Are-Equal-2, then he does not risk a grave wrong by eating meat – the badness of eating meat is the same size on the All-Animals-Are-Equal-2 view as it is on the Ordinary Morality view, and it remains plausible that the prudential reason in favour of eating meat, on the Ordinary Morality view, outweighs the reasons against eating meat on both the Ordinary Morality view and the All-Animals-Are-Equal-2 view.

To illustrate, suppose (again very simplistically) that the prudential reason is 0.01 units in favour of chicken and 0.001 in favour of vegetarian; the reason against eating animals is 1 unit against chicken, not at all against vegetarian. The ordinary morality view regards units of prudential reason as 10,000 times as valuable as the units of moral reason not to eat animals. Thus according to ordinary morality the value of eating chicken, for example, is given as $0.01 \times 10,000 - 1 = 99$. This gives us the following two tables.

If Harry has credence in All-Animals-Are-Equal-1 then it's clear that the moral risk of eating chicken is grave and that, unless Harry's credence in All-Animals-Are-Equal-1 were tiny, it would be inappropriate to eat chicken.

Option	Ordinary Morality	All-Animals-Are-Equal-1
Eat chicken	99	- 9,900
Eat vegetarian	10	10

²³Note that we need to include three outcomes in this table because we are interested in making intertheoretic comparisons of choice-worthiness *differences* rather than intertheoretic comparisons of choice-worthiness *levels*.

In contrast, if Harry has credence in All-Animals-Are-Equal-2, then the potential moral downside of eating chicken is much smaller. Indeed, the biggest potential loss of value is to fail to eat chicken if ordinary morality is correct. Harry would need to have a very low credence in ordinary morality in order for eating vegetarian to be the appropriate option.

Option	Ordinary Morality	All-Animals-Are-Equal-2
Eat chicken	99	– 0.99
Eat vegetarian	10	0.001

Because there are two distinct and seemingly natural ways of making the intertheoretic comparison, we again see that the moral uncertainty based argument for vegetarianism doesn't straightforwardly go through. We need to make a controversial decision about which of these two ways of making the intertheoretic comparison is correct.

A similar issue affects the moral uncertainty argument against abortion. As we noted above, we cannot say that there's no serious moral downside to keeping the child, because having a child costs resources that could be used to save lives. This argument becomes stronger when we consider the issue of intertheoretic comparisons.

Let us assume that Isobel has some credence in the view that there's no morally relevant distinction between acts and omissions. Again, there are two distinct but natural ways of doing the intertheoretic comparison. Let us suppose that Ordinary Morality regards a killing as 1,000 times as bad as a letting die. In which case, we can represent the two ways of normalizing the view that rejects the acts/omissions distinction as follows:

Option	Ordinary Morality	No-Acts/Omissions-1	No-Acts/Omissions-2
Kill 1 person	– 1,000	– 1,000	– 1
Let 1 person die	– 1	– 1,000	– 1
No change	0	0	0

On No-Acts/Omissions-1, letting die is far worse than Ordinary Morality supposes; it's as wrong as killing. On No-Acts/Omissions-2, killing is much less bad than Ordinary Morality supposes; it's merely as wrong as letting die.

If Isobel only has some credence in No-Acts/Omissions-2, then her credence in the idea that there is no acts/omissions distinction is not going to have a big impact on the appropriateness ordering of her options. If, in contrast, she has some credence in No-Acts/Omissions-1, then the biggest moral consideration in her decision whether to have an abortion is not the potential killing of an innocent person, but the opportunity cost of the resources that she would spend on the child, which could be used to prevent the deaths of others.

Once again, therefore, one cannot claim that the implications of MEC follow straightforwardly whatever set of reasonable credences one has. In addition to making potentially controversial claims about what credences one ought to have across different moral views, in order to come to a conclusion about what moral uncertainty considerations entail in a particular case one also must often make potentially controversial

claims about what is the correct way of making intertheoretic comparisons across the views that the decision-maker has credence in.²⁴

Note that none of what I've said so far is an argument for the conclusion that vegetarianism or anti-abortion views *don't* follow from consideration of moral uncertainty. All I've argued is that invoking moral uncertainty alone is not sufficient to conclude that vegetarianism is right or that abortion is wrong. Instead, one must also invoke substantive and potentially controversial assumptions about what credences one ought to have across a wide array of moral views, and across different choices of intertheoretic comparisons.

Nor am I arguing that moral uncertainty does not have concrete implications for real-life decision-makers. Once a decision-maker has determined at least approximately what her credences across different theories and across different intertheoretic comparisons are or ought to be, maximizing expected choice-worthiness will recommend some courses of action as appropriate and others not. I strongly suspect the resulting recommendations will look quite different from the typical positions in debates on these issues, or from the view that one would come to if one simply followed one's favoured moral view.

In general, and very roughly speaking, I believe that maximizing expected choice-worthiness under moral uncertainty entails something similar to a value-pluralist consequentialism-plus-side-constraints view, with heavy emphasis on consequences that impact the long-run future of the human race.²⁵ How that exactly plays out depends on which normative views, and manners of making intertheoretic comparisons, one finds most plausible. If one is sympathetic to consequentialism, or to intertheoretic comparisons that favour consequentialism (such as No-Acts/Omissions-1), then maximizing expected choice-worthiness might entail something like the principle: 'Maximize the good except when doing so would obviously violate a possible side-constraint for limited gain.' If, on the other hand, one is sympathetic to non-consequentialism, or to intertheoretic comparisons that favour non-consequentialism (such as No-Acts/Omissions-2), then maximizing expected choice-worthiness might entail something like the principle: 'As much as possible, avoid violating possible side-constraints, but if you are confident that you are not going to violate a side-constraint, then maximize the good.'

V. Conclusion

In this article, I've argued that the moral uncertainty based arguments that philosophers have given in the literature for the rightness of vegetarianism and the wrongness of abortion are too simple. The precise implications of maximizing expected choice-worthiness under moral uncertainty depend on potentially controversial assumptions

²⁴One might claim that (i) one ought to have credence in both possible normalizations and that (ii) given this, the theory with the higher-stakes normalization will still be the primary determiner of different options' expected choice-worthiness. I find this plausible to some extent, but believe it still depends on what exactly one's credences are; if one has a very small credence in the high-stakes normalization, then one might worry that one is entering 'fanaticism' territory if one thinks that the recommendation of MEC in this instance is correct.

²⁵For the argument why total view consequentialists should care almost exclusively about impacts on the long-run future of the human race, and why, on their view, there is a truly vast amount of value at stake, see Nick Bostrom, 'Astronomical Waste: The Opportunity Cost of Delayed Technological Development', *Utilitas* 15 (2003), pp. 308–14.

about what credences one ought to have across different moral views, and about how to make intertheoretic comparisons across theories.

I do believe, however, that consideration of moral uncertainty should have major impacts for how practical ethics is conducted. Currently, a central focus of practical ethicists is on determining what the most plausible view on a given issue is, by arguing in favour of that view, or by arguing against competing views. If moral uncertainty were taken into account, then an additional vital activity for practical ethicists to engage in, before any recommendations about how to act were made, would be to consider the implications of a variety of different moral views on this issue, to argue for what credences to assign to those views and for what the most plausible intertheoretic comparisons are, and then to work out which options have highest expected choice-worthiness. It would be surprising if the conclusions of this were the same as those that practical ethicists typically draw.²⁶

²⁶For their generous comments and discussion, I would like to thank Frank Artzenius, Amanda Askill, John Broome, Krister Bykvist, Toby Ord, Peter Singer, Christian Tarsney, Ralph Wedgwood, and audiences at the University of Oxford, The London School of Economics, the Princeton Center for Human Values Ira W. DeCamp Bioethics Seminar, and the Uehiro Centre Applied Ethics Discussion Group.