Journal of Risk Research

ISSN: 1366-9877 (Print) 1466-4461 (Online) Journal homepage: www.tandfonline.com/journals/rjrr20

Acceptable risk under moral uncertainty

Ibo van de Poel

To cite this article: Ibo van de Poel (14 Nov 2025): Acceptable risk under moral uncertainty, Journal of Risk Research, DOI: 10.1080/13669877.2025.2584028

To link to this article: https://doi.org/10.1080/13669877.2025.2584028

JOURNAL OF RISK RESEARCH
https://doi.org/10.1080/13669877.2025.2584028

Acceptable risk under moral uncertainty

Ibo van de Poel

Department Values, Technology & Innovation, TU Delft, Delft, the Netherlands

ABSTRACT

While there is a considerable literature on acceptable risk, which considers questions like under what condition is it acceptable to impose risks on others, this literature has neglected the issue of moral uncertainty. Moral uncertainty is uncertainty about what moral theory is right. I explore whether, and if so how, moral uncertainty is relevant for (moral) decisions about acceptable risk. To do so, I first consider whether a popular decision procedure to deal with moral uncertainty, Maximizing Expected Choiceworthiness (MEC) is an appropriate procedure for accounting for moral uncertainty in decisions about acceptable risk. I conclude it is not, due to reasons of intertheoretic incomparability and lack of justifiability of the resulting decisions. I then consider several alternatives – such as ignoring moral uncertainty, following the decision-maker's favorite moral theory and decision rules that avoid intertheoretical tradeoffs – that also turn out to be wanting. Next, I argue that the problems of MEC and alternative decision procedures go back to fundamental assumptions about how to deal with moral uncertainty. If these assumptions are lifted, the road is open to a more deliberative and coherentist approach to decision-making about acceptable risk under moral uncertainty that avoids the discussed problems.

ARTICLE HISTORY
Received 21 June 2025
Accepted 28 October 2025

KEYWORDS
Moral uncertainty; acceptable risk; descriptive uncertainty; moral theory; risk

1. Introduction

Decisions about acceptable risk often need to be made not just under descriptive uncertainty, but also under moral uncertainty. Consider, for example, the following situation. During the COVID-19 pandemic, governments had to make decisions about how to deal with the risks of the pandemic. At some point in time, there was a fear that the pandemic would spread so quickly that hospitals would get overcrowded. To decrease the transmission rate of the virus, options like lockdowns were considered: closing down shops and schools and other public places to avoid or at least reduce transmission of the virus. Lockdowns are drastic message in liberal democracies as they curtail basic freedoms of citizens. Moreover, they have economic consequences and may affect the mental health of, for example, youngsters. Apart from (presumably) having positive effects, lockdowns also introduce all kinds of harms and risks, which may make them unacceptable.

Such decisions about lockdowns were made not just under descriptive uncertainty but also under moral uncertainty. They involved descriptive uncertainty, because it was not known how effective they would be and whether they would indeed prevent the overcrowding of hospitals.

CONTACT Ibo van de Poel i.r.vandepoel@tudelft.nl Department Values, Technology & Innovation, TU Delft, Jaffalaan 5, 2628 BX Delft, the Netherlands

It was also not known, for example, what the exact effects on the economy and mental health would be. But such decisions also involve moral uncertainty. For example, should the decision be based on the greatest good for the greatest number or should governments take individual rights and liberties into account, and if so how should they weigh them? How should economic versus the health effects be weighted? And how should they weigh the interest and rights of different parts of the population?

There is a considerable literature on the ethics and philosophy of risk (e.g. Asveld & Roeser, 2009; Hansson, 2003, 2023; Hayenhjelm & Wolff, 2012; Roeser et al. 2012; Shrader-Frechette, 1991), and an already older literature on acceptable risk (e.g. Fischhoff et al. 1981; Glickman & Gough, 1990; Otway & Von Winterfeldt, 1982; Slovic, 2000). However, both bodies of literature have hardly discussed questions of acceptable risk under moral uncertainty.[1] My aim in this contribution is to consider whether, and if so how, moral uncertainty matters for questions about acceptable risk. For the purpose of this article, I will understand risk as the probability of harm of an unwanted event.[2]

In exploring how moral uncertainty should be accounted for in decisions about acceptable risk, I will build on the growing literature on moral uncertainty in philosophy (e.g. Lockhart, 2000; MacAskill et al. 2020; Ross, 2006; Sepielli, 2009). In particular, I will consider one popular decision procedure that has been proposed in this literature, namely Maximizing Expected Choiceworthiness (MEC) to see whether it is an appropriate and feasible way to deal with questions of acceptable risk under moral uncertainty. I will conclude that it is not, because it assumes intertheoretical comparability between moral theories, but also because it bars justifying decisions about acceptable risks in terms of underlying (moral) reasons, while such justification seems essential, particularly for public decisions that may affect large parts of the population, like the mentioned decision about lock-downs during a pandemic.

Next, I will consider several alternatives to MEC, namely ignoring mortal uncertainty altogether, basing the decision on the moral theory or ethical framework the decision maker has most credence in, or alternative decision rules that avoid intertheoretical trade-offs. I will argue that all these alternatives are also beset with problems. Therefore, I will propose to take a step back in order to question an assumption that is common to MEC and alternative decision rules for dealing with moral uncertainty. This is the assumption that in order to account for moral uncertainty, we first need to determine how each individual moral theory or ethical framework orders the options (taking into account descriptive uncertainty) and then add moral uncertainty to the picture. I will argue that this assumption is problematic for several reasons and that if we drop the assumption the way is open to a more pluralistic approach that allows taking into account moral uncertainty as one of the considerations in decisions about acceptable risk.

Such an approach not only allows for decisions about acceptable risk for which a justification in terms of underlying (moral) reasons can be given but also overcomes some other problems with MEC that have been discussed in the literature. The approach I propose means that under conditions of moral uncertainty, decision-makers should not simply follow one of the existing ethical frameworks for acceptable risk, but should also account for moral uncertainty. I suggest that this might be done through a coherentist approach (that aims for example for reflective equilibrium) rather than through the application of a decision rule to account for moral uncertainty.

2. Acceptable risk

Although descriptive uncertainty is abundant in the real world, most traditional moral theories have ignored it (Hansson, 2003). That is to say, for most traditional moral theories it is unclear, or at least open to interpretation, how they should be applied under conditions of descriptive uncertainty. How should a Kantian come to a moral judgement when there is 1% probability that someone is used as a means to an end? Is that equally unacceptable as for sure using

someone as a means to an end? Is there a probability limit below which we should not worry or is any probability, whatever small, unacceptable?

Utilitarianism is an exception to the rule that traditional moral theories have tended to ignore descriptive uncertainty. It can, and has been, combined with the notion of expected utility, for example based on the axiomatization of expected utility by Von Neumann and Morgenstern (1944), or another axiomatization. The idea is that we should maximize the probability-weighted average of the utility of the possible outcomes. This approach is known as maximizing expected utility.

It has been suggested that this approach can be generalized to maximizing expected value, where value can also be non-utilitarian, so accounting for non-utilitarian moral theories (cf. Hedden, 2015; Nissan-Rozen, 2015; Oddie & Milne, 1991). Nevertheless, it is questionable whether all moral theories can be projected on a uniform value function, as would be required for maximizing expected value (Hedden, 2015). Moreover, it is also questionable whether all moral theories will account for descriptive uncertainty in a similar way (Hansson, 2003; Lazar & Graham, 2021; Tenenbaum, 2017). Still, the ideal of maximizing expected value seems to have motivated the idea that moral uncertainty should be accounted for through Maximizing Expected Choiceworthiness (MEC) (e.g. MacAskill et al. 2020), which I will discuss in the next section.

In studies on acceptable risk, the idea of maximizing expected utility or value has been further operationalized through what is known as risk-benefit analysis or cost-benefit analysis, which aim at selecting the alternative that has the maximum expected (monetary) value (Shrader-Frechette, 1991; Sunstein, 2005). Risk-benefit analysis has been criticized from a philosophical and ethical point of view (Hansson, 2007b). Moreover, studies on risk perception have shown - since at least the 1980s - that many other factors play a role in what risks people are willing to accept (e.g. Fischhoff et al. 1981; Glickman & Gough, 1990; Slovic, 2000). Examples of such factors are:

- The voluntariness of the risk;
- The distribution of the risks and benefits;
- The maximum negative outcomes;
- The controllability and repairability of the possible negative outcomes;
- The familiarity with, and proximity and visibility of certain risks.

For some of these factors, it might perhaps be argued that taking them into account in decisions about acceptable risk is irrational (or even immoral). This applies, for example, to the familiarity, proximity and visibility of risks. However, for some of the other factors, it is clear that they represent legitimate ethical concerns, like, for example, the voluntariness of risks and the distribution of risks and benefits.

There is now a considerable literature on acceptable risk. In general, the focus in the literature has moved from questions about whether individual risks are acceptable, to questions about the acceptability of risky activities (that have also benefits), to more recently, an emphasis on under what conditions it is acceptable to take risks, or to impose risks on others (Hansson, 2023; Hayenhjelm & Wolff, 2012).

Different moral principles or ethical frameworks have been proposed in the literature to judge the acceptability of risks and of risking, i.e. the activity of introducing risks to oneself (risk taking) or to others (risk imposition), including the following:[3]

- 'Exposure of a person to a risk is acceptable if and only if the total benefits that the exposure gives rise to outweigh the total risks, measured as the probability-weighted disutility of outcomes'. (the traditional criterion in risk analysis according to Hansson (2003: 306)).

- Risky activities are acceptable if their probability-weighted benefits outweigh the probability-weighed harms and there is no other activity (that can attain the same ends) with a higher expected value.[4] (maximizing expected value applied to acceptable risk)
- Exposure of a person to a risky activity is acceptable if they have voluntarily consented to the risk after been fully informed about the risks and benefits of that activity. This principle is often applied in (individual) medical contexts (Beauchamp & Childress, 2013), but has also been proposed for technological risks (Martin & Schinzinger, 2005)).[5]
- 'Exposure of a person to a risk is acceptable if and only if this exposure is part of an equitable social system of risk-taking that works to her advantage'. (Hansson, 2003: 305).
- Imposition of a risk is acceptable to the extent it does not harm the risk-bearer (where harm may be understood in different ways) (e.g., Finkelstein, 2003; Placani, 2017).
- Imposition of a risk is acceptable to the extent that it does not reduce the overall freedom of the risk-bearer (Ferretti, 2016).
- Imposition of a risk is acceptable to the extent it does not infringe on the autonomy (and the available options) of the risk bearer (Oberdiek, 2017).
- Imposition of a risk is acceptable to the extent that the negative effects of the risk, when it materializes can (either) be reversed, repaired or compensated (Hayenhjelm, 2018).
- Imposition of a risk is acceptable to the extent that the risk imposition is non-dominating, i.e., if it does not allow the risk imposer to exercise their arbitrary power over the range of safe options that the risk-bearer are in a position to choose from (Maheshwari & Nyholm, 2022).

3. Moral uncertainty and MEC (maximizing expected choiceworthiness)

Given that there are different ethical frameworks for acceptable risk, as we have seen in the previous section, and given that these frameworks might lead to different decisions about what risk are acceptable, it is very well conceivable that a decision-maker faced with a decision about acceptable risk is morally uncertain about what ethical framework to apply.

In the philosophical literature, moral uncertainty is usually defined as uncertainty about the moral facts and it needs to be distinguished from descriptive uncertainty which is uncertainty about descriptive facts. Oftentimes, moral uncertainty is understood as uncertainty about what moral theory is right. In our case, we have different ethical frameworks for acceptable risk, and I will understand moral uncertainty as uncertainty about which of these ethical frameworks is right and should be applied to make decisions about acceptable risk.

Different approaches for decision-making under moral uncertainty have been proposed in the literature (Gustafsson & Torpman, 2014; Kernohan, 2021; Lockhart, 2000; MacAskill et al. 2020; Riedener, 2020; Sepielli, 2009; Tarsney, 2019). For now, I focus on the idea that under conditions of moral uncertainty, we should aim to maximize expected choiceworthiness (MEC). I will focus in particular on the proposal made by MacAskill et al. (2020). They understand choiceworthiness as representing 'the strength of the reasons for choosing an option' (MacAskill et al. 2020: 4) and they think that decision makers should maximize choiceworthiness across the moral theories in which the decision maker has some credence.

The reason to focus on MEC is that this approach is currently the best elaborated approach to deal with moral uncertainty, and the proposal by MacAskill et al. (2020) is probably currently the most sophisticated and detailed version of MEC. Although I will criticize MEC in the next section and argue that it is not an appropriate approach for making decisions about acceptable risk under moral uncertainty, I nevertheless think that MEC is probably the best established and developed proposal for dealing with moral uncertainty. Moreover MEC is interesting because it builds on similar ideas and aims as maximizing expected utility or value for situations of descriptive uncertainty. In particular, it builds on the idea that we have good reasons to treat

descriptive and moral uncertainty in a similar fashion in moral decision-making (MacAskill et al. 2020: 47-50). For all these reasons, MEC is a useful starting point for our investigation

MEC assumes that the decision-maker has *credences* (degrees of subjective belief) in various moral theories. Each of these moral theories might order the available options differently. The idea of MEC is that in making a decision under moral uncertainty, we should not only take these credences into account but also how good (or bad) the various options are according to the various moral theories, so allowing for trade-offs (MacAskill et al. 2020: 44-47).

As an illustration, consider how MEC would help to make a decision under moral uncertainty for the COVID-19 example which I presented in the introduction. For the sake of simplicity, I will be assuming that the decision-maker considers three options, namely 1) no lock-down, 2) a partial lock-down, and 3) a full lock-down. Each of these options has a number of morally relevant consequences, like hospital occupancy rates, numbers of fatalities, economic effects, effects on the mental health of youngsters, infringements on the liberty of people, etc. For the sake of simplicity, I am assuming that the decision maker is descriptively certain about these consequences, and is only morally uncertain how to morally evaluate these consequences. Now, we might further assume that the decision-maker considers the following three ethical theories as having some credence:

1. A utilitarian theory that tells to select the option that will lead to the greatest benefit for the greatest number (in the long run);
2. A social justice theory that gives special weight to protecting the most vulnerable (e.g., elderly people with a disease who may be particularly vulnerable to COVID-19);
3. A libertarian theory that postulates that infringements on people's rights and freedoms should be minimalized, and are only justified if there is an imminent danger.

Table 1 indicates how MEC might help to make a decision in this case. For each moral theory, the decision maker has to give their credence in that theory, and the decision maker needs to determine the choiceworthiness of each option under each moral theory. Now, in most cases moral theories dot not directly provide choiceworthiness scores, that is to say they usually provide an ordering of the options in terms of choiceworthiness, but do not provide choiceworthiness numbers. Nevertheless MEC assumes that under certain conditions, it is possible to construe choiceworthiness scores that are intertheoretically comparable. A hypothetical example is shown in Table 1. The decision-maker can now calculate the expected choiceworthiness of each option using the following formula:

$$c_m = \sum_{i=1}^{n} p_i c_{m,i}$$

with c_m the expected choiceworthiness of the m^{th} option, p_i the credence in the i^{th} moral theory, $c_{m,i}$ the choiceworthiness of the m^{th} option according to the i^{th} moral theory and n the number of relevant moral theories.

In cases where it is not possible to construe numerical choiceworthiness scores on basis of the relevant moral theories, MacAskill et al. (2020) suggest to use the Borda Count, which only requires that the theories provide an ordinal ordering of the options. The Borda Count awards

Table 1. MEC applied to the lockdown example; p stands for the credence the decision maker has in an ethical framework, c for the choiceworthiness of an option according to an ethical framework.

	Utilitarianism		Social justice		Libertarian		Expected choiceworthiness
Option	p	c	p	c	p	c	
No lockdown	0.3	30	0.5	10	0.2	60	26
Partial lockdown	0.3	50	0.5	20	0.2	5	26
Full lockdown	0.3	20	0.5	50	0.2	0	31

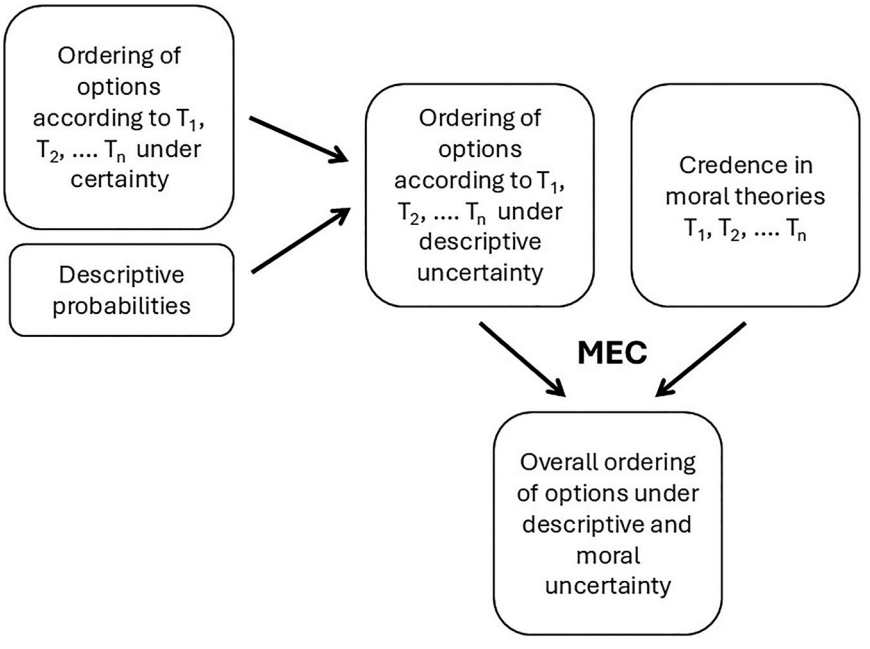

Figure 1. How MEC models moral decision-making under descriptive and moral uncertainty.

Table 2. MEC with the Borda Count applied to the lockdown example; p stands for the credence the decision-maker has in an ethical framework, BC for the Borda Count of an option according to an ethical framework.

	Utilitarianism		Social justice		Libertarian		Expected
Option	p	BC	p	BC	p	BC	choiceworthiness
No lockdown	0.3	2	0.5	1	0.2	3	1.7
Partial lockdown	0.3	3	0.5	2	0.2	2	2.3
Full lockdown	0.3	1	0.5	3	0.2	1	2

a score of 1 to the option that is ordered lowest according to a moral theory, a score of 2 to the option above, and so on. The result for our example is shown in Table 2.

4. Criticisms of MEC

A main criticism of MEC and similar theories for dealing with moral uncertainty is that they assume some form of intertheoretical comparability (Gracely, 1996; Gustafsson & Torpman, 2014; Hedden, 2015; Nissan-Rozen, 2015). The reason why MEC is sensitive to such a criticism is that the procedure is based on trade-offs between different moral theories, which requires some form of intertheoretical comparability.

MacAskill et al. (2020) present a positive argument for intertheoretical comparability. They suggest that there may be a universal (ratio) scale for choiceworthiness grounded in (moral) reality. But, apart from meta-ethical and metaphysical doubts one might have about this claim, even if such a scale would exist, it is unclear how the ordering of options given by different moral theories can be mapped on this scale. Without such a mapping, the existence of a uniform ratio scale for choiceworthiness is of little practical help in achieving intertheoretical comparability. Consider Table 1, the existence of a universal scale of choiceworthiness suggests that each moral theory assigns a choiceworthiness to each option, but this does not help to determine what choiceworthiness numbers to fill in in the table.

It might seem that if we use the Borda Count, as is done in Table 2, the problem of intertheoretical incomparability is avoided, but that impression is wrong. In fact the calculation of an overall expected choiceworthiness in Table 2 means that the Borda scores according to different theories are interpreted as being intertheoretically comparable. Without this assumption, it is not possible to calculate the overall expected choiceworthiness, and no choice can be made.

While I believe that intertheoretical incomparability is a problem for MEC, I will focus here on another criticism, namely the problem that MEC does not allow for a justification of decisions about acceptable risk in an appropriate way. This point is – as far as I am aware – new and has not been articulated in detail before. Although in driving home the point, I will focus on the case of public decisions about acceptable risk, the argument is more general and raises doubts also about the application of MEC in other contexts as well.

4.1. The need for justification

Decisions about acceptable risk need to be justifiable. This is particularly true if such decisions imply an imposition of risks on others. Roughly, justifiable here means that the decision maker can explicate the (moral) reasons for the decision.

The notion of justifiability I have in mind here does not imply that the decision maker has made a decision that is also right from a more objective point of view. Such (objectively) right decisions may be what I will call strongly justified, e.g. if they meet something like the following condition:

> **Strong justification**: a moral decision is strongly justified if it is (objectively) the right decision and the decision maker can explicate the right reasons for the decision.

A strongly justified decision is thus both (objectively) right and the decision maker has access to – and can explicate or cite - the appropriate (moral) reasons that justify the decision.

However, under conditions of (descriptive and moral) uncertainty we cannot expect decision makers to make the (objectively) right decisions as it its – by definition – unknown what the (objectively) right decision is. Still, it seems reasonable to require that the decision maker can justify the decision in a weaker sense. I propose to understand this weaker sense as follows:

> **Weak justifiability**: the decision maker can subjectively justify the decision by citing the (moral) reasons that underlie the decision (according to the decision maker) and can indicate why they believe that on the balance of reasons the decision is right.

The justification here is subjective because it is relative to the belief states of the decision maker. Importantly, weakly justifiable decisions require the decision maker not just to describe the process (or procedure) they went through to make a decision but also to explicate the underlying (moral) reasons, or at least what the decision-maker believes to be the underlying reasons that justify the decision.

We typically require that decisions about acceptable risks are weakly justifiable. How does this play out for moral decisions about acceptable risk under moral uncertainty? First, it is worth noting that if decision makers base their decisions about acceptable risk on one of the normative frameworks for acceptable risk mentioned at the end of section 2, they would typically be able to weakly justify their decisions. Such justifications may, for example, read as follows:

- These risks are acceptable because overall society is better off by accepting them.
- These risk are acceptable because they do not infringe on people's rights.
- These risks are acceptable because people have consented to them.
- These risks are acceptable because they do not imply a dominating risk imposition.
- These risks are acceptable because they are a part of arrangements that work to the benefit of all and that protect the basic rights of all.

However, when the decision maker wants to take moral uncertainty into account and makes a decision based on MEC, a (weak) justification seems no longer available. The reason is that MEC considers different orderings of the options, according to different ethical frameworks, but that it is unclear how the *reasons* that inform these frameworks add up or are combined in the overall decision.

So suppose the government of some country had used MEC to decide about lock-downs, and decided for a full lockdown, e.g. based on the numbers in Table 1. Asked for a justification about their decision, they might say something along the following lines: 'This calculation shows that a full lockdown has the highest expected choiceworthiness. You can calculate it for yourself'. But that seems quite technocratic, and not very convincing and satisfying as a justification of such an important decision. This would seem an important objection against MEC.

How can a proponent of MEC reply? I think two replies are possible: 1) MEC allows for a procedural (second-order) justification of the decision made and that is perfectly acceptable as we also accept such justification in other contexts (like democratic decision-making), 2) MEC makes it more likely that we make the morally right decision and that is more important than making a justifiable decision. I will briefly consider these two replies.

First, it is indeed true that MEC can offer a procedural justification of the decision, but it remains questionable whether that is good enough. The fact that we (rightly) accept procedural justification for democratic decision-making does not mean that we should also accept it for other decision-making procedures. Moreover, MEC does not seem a viable and acceptable alternative to current democratic decision-making because it is - as currently proposed - a technocratic procedure that lacks democratic legitimation, as it assumes *one* decision-maker that has different credences in various moral theories. Perhaps, it might be possible to propose new democratic procedures based on MEC, but that seems a new project and not MEC as it is currently proposed.

With respect to the second argument, it might be questioned whether MEC indeed makes it more likely to come to a morally right decision. Of course, the proponents of MEC believe so, but not everyone agrees, for example due to problems of intertheoretical incomparability. Second, even if MEC makes it more likely to make right decisions, it can be questioned whether that is worth the price of not be able to justify such decisions. Perhaps it would be worth this price if MEC would be the only way to take moral uncertainty seriously, but it is not. There are alternative ways to take moral uncertainty seriously, as we will see below, and they do not face the problem of a lack of (weak) justifiability.

5. Alternatives

In this section, I will discuss three alternatives for MEC for decision-making about acceptable risk under moral uncertainty: 1) ignoring moral uncertainty altogether, 2) My Favorite Theory and 3) other decision procedures that avoid intertheoretical trade-offs. We will see that the last two has some advantages over MEC in terms of the two objections against MEC I have discussed. Nevertheless, as will see, they also have some problems of their own.

5.1. Ignoring moral uncertainty

Some authors have argued that we should ignore moral uncertainty altogether (Harman, 2015; Hedden, 2015; Weatherson, 2014). One argument for such a position is the impossibility of intertheoretical comparisons (e.g. Hedden, 2015). But, although this is an argument against MEC, it is not an argument against moral uncertainty because there are ways to deal with moral uncertainty that do not require intertheoretical comparison, like for example My Favorite Theory (which I will discuss below).

Two other arguments against moral uncertainty stand out. First, the argument that moral uncertainty is different from descriptive uncertainty and should therefore be treated differently, and – second – the argument that moral uncertainty leads to improper moral motivations,

which in the literature are called 'fetish motivations'. I will discuss both briefly and argue that neither is a convincing argument against taking moral uncertainty seriously.

The first argument is neatly summarized by Hedden (2015, 102) as follows: 'The motivations for wanting a theory of what to do given descriptive uncertainty do not carry over to normative uncertainty. Descriptive facts may be inaccessible even in principle, and (non-culpable) ignorance of them excuses one from blame, but normative facts are in principle accessible, and ignorance of them arguably is no excuse'.

The problem with this argument is that even if it were true (of which I am less than certain) that moral or normative uncertainty is relevantly different from descriptive uncertainty, we still have good reasons to think about how to make decisions under moral uncertainty. After all, even if 'normative facts are in principle accessible', in an epistemological and practical sense they may very well be inaccessible, even for a serious moral agent. Perhaps we should account for moral uncertainty in a different way than for descriptive uncertainty, but that is not an argument for ignoring it.

The other part of Hedden's argument refers to the idea that while descriptive ignorance may sometimes be excusing, moral uncertainty never is. This goes back to a similar argument by Harman (2015). However, such an argument relies on the problematic assumption that moral ignorance is never an excuse, which seems an intuition that is not widely shared (cf. MacAskill et al. 2020: 35). Moreover, it remains unclear what an agent who is genuinely morally uncertain and has consciously tried to reduce that uncertainty as much as possible is supposed to do more to avoid blame for an action that turns out to be (objectively) wrong. Again, there seems to be good reasons to consider moral uncertainty rather than just ignore it.

The second main argument against moral uncertainty is that it leads to what has been called a 'fetishist' moral motivation (Harman, 2015; Weatherson, 2014). The underlying idea here is that if moral action is solely based on formal considerations with respect to goodness (or rightness) or what has been called goodness *de dicto* (rather than *de re*) that it is fetishist (Smith, 1995). The idea is that there is something morally wrong with such motivation, which may be related to the fact that it does not track the right moral reasons, a point that seems related to the point that I have made in the previous section that MEC leads to decisions that are not (weakly) justifiable.

However, this does not seem to be a good argument for ignoring moral uncertainty. First, if a form of moral fetishism is the only way to take moral uncertainty seriously, it might be a price worth paying (Sepielli, 2016), Second, there are other ways for dealing with moral uncertainty – like the alternatives I will discuss below - that do not lead to a fetishist motivation, as we will see, and that take moral uncertainty (more) seriously.

5.2. My Favorite Theory

The basic idea of My Favorite Theory is quite simple: the decision maker should choose the moral theory in which they have most credence and decide on basis of that theory. In the exemplary case of lock-downs, this would mean that the decision maker only considers the ethical framework of 'social justice' and chooses for a full lockdown (see Table 1).

This immediately leads to what I think is the main objection to this approach: while My Favorite Theory may formally acknowledge moral uncertainty, in the actual decision-making moral uncertainty is basically ignored. Nevertheless, My Favorite Theory is attractive for several reasons. It avoids the problem of intertheoretical incomparability (because it only considers one moral theory), and it allows for (weak) justifiability and avoids fetishist motivation (because it only concerns the moral reasons based on one moral theory).

My Favorite Theory has been defended against several possible objections by Gustafsson and Torpman (2014). I will discuss two of these objections briefly because they well illustrate the problems that My Favorite Theory gives raise to.

First, My Favorite Theory may give counterintuitive results if the decision maker has a slightly lower credence in a moral theory that judges the options quite differently than the decision maker's favorite theory. Gustafsson and Torpman (2014) consider the following situation (Table 3):

In this case, it would intuitively seem better to choose A, even if theory T_1 in which the decision-maker has the highest credence orders B over A. The reason is that so a grave moral wrong may be avoided. Gustafsson and Torpman (2014) reject this objection because it relies on intertheoretical comparison. While that is true and while with Gustafsson and Torpman (2014) I am sceptical about full intertheoretical comparability, it should be noted that this example requires a relatively modest form of intertheoretical comparability.

The second objection to My Favorite Theory has to do with the issue of theory individuation. If the decision maker is to decide in which theory they have most credence, they first need to individuate moral theories. Gustafsson and Torpman (2014) propose the following individuation account for the context of decision-making under moral uncertainty: 'Regard moral theories *T* and *T'* as versions of the same moral theory if and only if you are certain that you will never face a situation where *T* and *T'* yield different prescriptions' (Gustafsson & Torpman, 2014: 171).

However, as they are well aware, this principle may lead to counterintuitive results. Consider the case that the decision maker has 0.2 credence in utilitarianism (which orders A over B) and 0.8 credence equally divided over 10 deontological theories (i.e. 0.08 credence in each) which all order B over A. It would seem intuitive that the decision maker should prefer B over A but My Favorite Theory in combination with the proposed principle for theory individuation would lead to the choice of A. Gustafsson and Torpman (2014) and Gustafsson (2022) seem to believe that we should accept this result unless we can come up with a better proposal for theory individuation but that is not really an argument against the intuitive appeal of the counterexample.

5.3. Alternative decision procedures that avoid trade-offs

One of the reasons why MEC runs into issues of intertheoretical incomparability and lack of (weak) justifiability is that it makes credence-weighted trade-offs between moral theories.[6] This suggests that if we can avoid such trade-offs we might also avoid these problems as well. Therefore, I will now look at two alternative decision rules that avoid (direct) trade-offs, namely My Favorite Option and the maximin rule.

Gustafsson and Torpman (2014: 165) characterize My Favorite Option as follows (for a similar principle, see Lockhart (2000: 26)):

> "An option *x* is a morally conscientious choice for *P* … if and only if *P* … has at least as high credence in *x* being right as in every other option."

One way in which this can be further operationalized is by making pair-wise comparisons between the options, and choosing the so-called Condorcet winner, i.e. the option that beats all other options in a pair-wise comparison (or at least does not lose). Table 4 illustrates this for our example.

Table 3. An example in which My favorite theory gives implausible results; p stands for the credence the decision-maker has in an ethical theory.

	Theory T_1 ($p = 0.51$)	Theory T_2 ($p = 0.49$)
Option A	slightly nasty	saintly
Option B	merely okay	terrible

However, the problem is that there is not always such a Condorcet winner because sometimes there are cycles, in which option A is ordered over B, B over C, and C over A, so that no choice is possible.

An alternative way of applying My Favorite Option is to interpret the outcomes of moral theories not as providing an ordering of options but simply in terms of classifying options as permissible or non-permissible. Under moral uncertainty, we can then select the option that is most likely to be permissible, see Table 5 for an example.

While this may be an attractive procedure under some circumstances, a problem is that, while for some moral theories – like deontological theories – it is quite natural to interpret them as classifying options as either permissible or non-permissible, for other moral theories - like utilitarianism – this is more artificial. One way to interpret utilitarianism is to say that only the option with the highest utility is permissible, but another (more plausible?) interpretation is that more options can be permissible, but then the question is where to draw the line between permissible and non-permissible options.

All in all, My Favorite Option has a number of advantages. The rule does not assume intertheoretical comparability and the resulting decision is (weakly) justifiable, because the decision maker can cite as reasons for the decision the reasons given by the theories that deem the decision right, while also acknowledging that there are reasons that over overridden, i.e. the reasons implied by those theories that deemed the decision wrong. For similar reasons, the rule does not seem sensitive to the problem of fetishist motivation. Still, the rule is sensitive to a number of other objections. If it is interpreted in terms of pairwise comparisons, it does not always provide an outcome. Alternatively, interpreted in terms of permissibility versus non-permissibility, it requires deciding where each moral theory draws this distinction which may be arbitrary and may no do justice to the gist of the underlying moral theory.

Another decision rule that does not require (direct) trade-offs is the maximin rule. This rule is aimed at avoiding the worst possible outcome by requiring the decision maker to choose the option of which the worst possible outcome is best compared to the worst possible outcome of other options. To do so, the decision-maker should for each option determine what the worst possible outcome is if each of the relevant ethical frameworks turns out to be right, and then - by comparing these worst possible outcomes - select the option with the best worst possible outcome. See Table 6 for an example.

Table 4. Pairwise comparison applied to the lockdown example; p stands for the credence the decision-maker has in an ethical framework, c for the option preferred by an ethical framework for each pairwise comparison.

	Utilitarianism		Social justice		Libertarian		
Pairwise comparison	p	c	p	c	p	c	Winner
No lockdown versus partial lockdown	0.3	Partial	0.5	Partial	0.2	No	Partial
Partial lockdown versus full lockdown	0.3	Partial	0.5	Full	0.2	Partial	Tied
No lockdown versus full lockdown	0.3	No	0.5	Full	0.2	No	Tied

Table 5. My Favorite option applied to the lockdown example for expected permissibility; p stands for the credence the decision maker has in an ethical framework, P means that the option is permissible according to an ethical framework, NP that the option is not permissible.

	Utilitarianism		Social justice		Libertarian		Expected permissibility
Option	p		p		p		
No lockdown	0.3	P	0.5	NP	0.2	P	0.5
Partial lockdown	0.3	P	0.5	NP	0.2	NP	0.3
Full lockdown	0.3	NP	0.5	P	0.2	NP	0.5

Table 6. Maximin applied to the lockdown example; p stands for the credence the decision-maker has in an ethical framework, c for the choiceworthiness of an option according to an ethical framework.

	Utilitarianism		Social justice		Libertarian		Minimal choiceworthinesss
Option	p	c	p	c	p	c	
No lockdown	0.3	30	0.5	10	0.2	60	10
Partial lockdown	0.3	50	0.5	20	0.2	5	5
Full lockdown	0.3	20	0.5	50	0.2	0	0

It would seem that this decision rule allows the decision maker to (weakly) justify the decision by pointing out the substantive reasons why it is important to avoid outcomes that are worse than the outcomes possible with the chosen option (such reasons may be suggested by the moral theory that selects other options as having worse possible outcomes than the one selected). However, the maximin rules assumes some intertheoretical comparability of moral theories; more specifically, it requires that the worst possible outcomes according to various moral theories are level comparable. It might be questionable whether such level comparability is achievable and meaningful.

The maximin rule has other disadvantages as well. One is that it neglects credences in moral theories. This may particularly problematic when a theory A in which the decision-maker has a low credence orders some option as worst, while other theories (in which the decision maker has a high credence) scores this option (very) high. For example, Table 6 shows that in this case the libertarian ethical framework is decisive in choosing for no lockdown, while the decision-maker has only a credence of 0.2 in this framework.

6. Restating the problem

I have argued that MEC is not an appropriate procedure for decision-making about acceptable risk due to problems of intertheoretical incomparability and lack of (weak) justifiability of decisions. I then considered alternatives that all have some advantages over MEC but are also beset with (other) problems. Moreover, the various procedures lead to different outcomes when applied to our example (see Table 7). Perhaps, this means that MEC is not so bad after all. But rather than concluding that, I will take a step back and will question some of the more fundamental assumptions underly MEC and other decision procedures that try to capture moral uncertainty.

The way that MEC, and other decision procedures - like My Favorite Theory and My Favorite Option - deal with moral uncertainty rests on two fundamental assumptions, which both are questionable, as I will argue below. The first assumption is that to account for moral uncertainty, we should first, for each moral theory, determine how it orders the options under descriptive uncertainty, and then, at the end of the pipeline so to speak, add moral uncertainty to the picture. The second assumption is that only one moral theory is right but the decision maker does not know which one and that we therefore should account for moral uncertainty in terms of credences in moral theories.

I will argue below that both assumptions are problematic and that when we lift them, this opens the ways to alternative, more adequate ways for dealing with moral uncertainty in decisions about acceptable risk.

Let me start with the first assumption. Figure 1 represents how decision procedures like MEC currently account for moral uncertainty. We have a range of moral theories T_1, T_2, …. T_n that each orders the options under descriptive certainty. We then add descriptive uncertainty to the picture resulting in each of the theories ordering the options under descriptive uncertainty. Next, we add moral uncertainty through credences in moral theories.

One problem with this picture is that different moral theories are likely to deal with descriptive uncertainty in different ways, as I already briefly pointed out in section 2 (see also Hansson,

Table 7. Different outcomes for the different decision rules for moral uncertainty when applied to the lockdown case.

Decision rule	Outcome
MEC	Full lockdown
MEC with Borda Count	Partial lockdown
My Favorite Theory	Full lockdown
My Favorite Option: pairwise comparison	Partial lockdown
My Favorite Option: expected permissibility	No lockdown or full lockdown
Maximin	No lockdown

2003; Lazar & Graham, 2021; Tenenbaum, 2017). Therefore, in section 3, I understood the problem of decision-making about acceptable risk under moral uncertainty in terms of different theories of acceptable risk R_1, R_2, …. R_n, in which each theory of acceptable risk already accounts for descriptive uncertainty (but not necessarily all in the same way). See Figure 2.

But Figure 2 is also problematic. It assumes that all ethical frameworks for acceptable risk should account for moral uncertainty in the same way, and that moral uncertainty only becomes relevant at the very end of the decision-making process. In other words, it presents the problem as a problem of mixing the outcomes of acceptable risk frameworks with moral uncertainty. This mixture problem is somewhat similar to the mixture problem described by Hansson (2003) for how (traditional) moral theories deal with descriptive uncertainty. His conclusion, and I think rightly so, is that it is better to understand this problem in another way, in which descriptive uncertainty is one of the factors.

A similar move can, and I think should, be made for the case of moral uncertainty. Rather than as conceiving of moral uncertainty as a factor to be added at the end of the process, it is relevant for deliberations about acceptable risk from the start, cf. Figure 3.

This brings me to the second problematic assumption in MEC and other decision rules for dealing with moral uncertainty: the idea that that we should account for moral uncertainty in terms of credences in moral theories because only one moral theory or ethical framework is right. This is problematic for two reasons. First, it is based on a too limited notion of moral uncertainty and, second, it takes moral theories to be mutually exclusive, rather than complementary.

Moral uncertainty is more than uncertainty about which moral theory is right. I can think of at least four reasons why not all moral uncertainty can be expressed in terms of credences in moral theories.

First, there can be moral uncertainty about how to apply a moral theory or how to exactly interpret a moral theory. However, proponents of MEC and other decision procedures for moral uncertainty might object that different applications or interpretations of a moral theory should simply be seen as alternative moral theories.[7] Although, this may lead to a proliferation of moral theories - which may be inconvenient from a practical point of view - in principle it seems possible to deal with moral uncertainty about the application or interpretation of moral theories in terms of competing moral theories.

Second, there might be - what might be called - moral framing uncertainty, that is uncertainty about how the problem of acceptable risk should be framed or understood in a specific case. For example does the case involve the imposition of risk on others? Can the risks when they materialize harm people? Does the risk imposition affect the autonomy of the risk bearers? It might be argued that decision-makers are often morally uncertain about the framing of acceptable risk problems and that different framings make different moral theories or ethical frameworks relevant. This moral framing uncertainty seems to be different from uncertainty about what ethical theory is right. It might, for example, manifest as moral uncertainty about what moral theory or ethical framework is *relevant* rather than as moral uncertainty about what moral theory is *right*. Nevertheless, proponents of MEC and other decision procedures might argue that framing is part of each moral theory and that, in as far as a moral theory allows

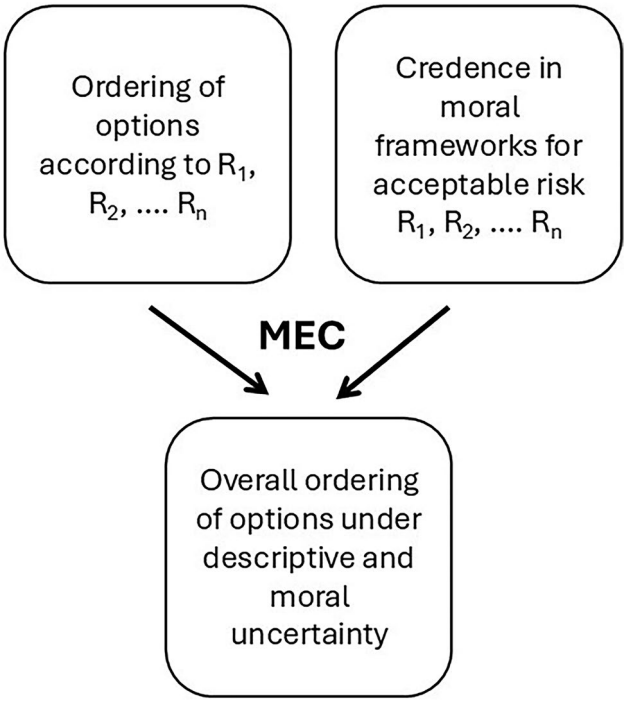

Figure 2. MEC applied to different moral frameworks for acceptable risks (which each take into account descriptive uncertainty).

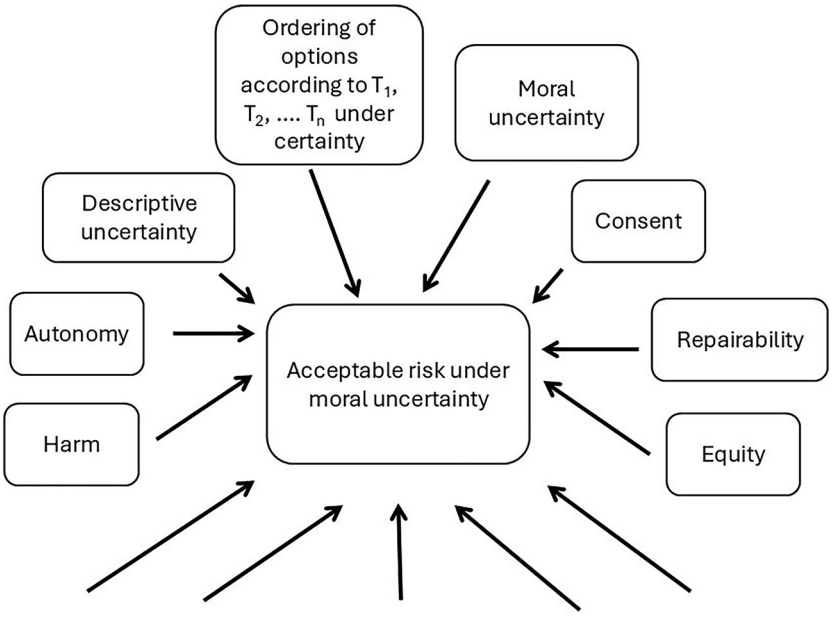

Figure 3. An alternative deliberative model for accounting for moral uncertainty in decisions about acceptable risk.

different framings, this – again - should simply be treated as different moral theories. However, this seems to ignore that moral framing uncertainty is not about credences in moral theories but rather about the relevance of such theories.

Third, moral uncertainty can give rise to new moral reasons that otherwise would not exist (Hicks, 2022). This can be moral reasons to reduce moral uncertainty before a decision is made, but also moral reasons to avoid certain risks, or reasons to take additional precautions, or to develop alternative options in order to be able to better deal with the moral uncertainty. Think, for example, of making more adaptive decisions, which can be adjusted when there is less moral (and descriptive) uncertainty (cf. Haasnoot et al. 2013; Hayenhjelm, 2018; Klinke & Renn, 2012).

Fourth, moral uncertainty can give rise to new moral norms or moral principles to better deal with that moral uncertainty. This point is related to the fact that different moral theories might not account for moral uncertainty in the same way. Particularly deontological moral theories might account for moral uncertainty not so much in terms of credences or probabilities, but rather in terms of (additional) moral norms and principles for decision-making under moral uncertainty (e.g. Lazar & Graham, 2021). In fact, this is already the case for descriptive uncertainty. When we consider the ethical frameworks that I mentioned in section 2 for dealing with acceptable risk, we can interpret these frameworks as articulating specific moral norms or principles for a particular kind of moral-decision making under descriptive uncertainty, namely situations in which new risks are introduced.

Not all moral uncertainty can therefore be expressed in credences in moral theories. More generally, a decision-theoretical framework might be too limiting for dealing with moral uncertainty (Rosenthal, 2021). Expressing moral uncertainty in credences is also dubious because it assumes that only one moral theory is true or right.[8] An alternative view sees moral theories not so much as competing but as complementary, that is to say, moral theories may be seen as articulating different moral reasons that may all be relevant for moral decisions about acceptable risk.

This alternative view assumes not so much a decision-theoretical approach to moral uncertainty but rather a deliberative approach, for example based on some form of coherentism, like a reflective equilibrium approach (Daniels, 1996; Rawls, 1999 [1971]). Such an approach fits well with Figure 3, in which moral uncertainty and descriptive uncertainty are additional concerns that feed into moral decisions about acceptable risk. Concerns that might give rise to new moral reasons or that may affect other moral reasons, and that also might require new moral norms and principles for decision-making about acceptable risk under moral uncertainty.

While such a deliberative and coherentist approach might require taking into account moral reasons articulated by different moral theories, it does not require intertheoretical comparability, at least not in the quantitative form that MEC does. Moreover, it enables for (weakly) justifiable decisions, as it requires the decision maker to articulate the relevant (moral) reasons as part of the deliberative process. For similar reasons, it avoids the problem of fetishist motivation, and – finally – it also allows treating descriptive and moral uncertainty differently, if and in as far as that is deemed desirable.

One way that such more deliberative and coherentist decision making about acceptable risk under moral uncertainty might take shape is through what has been called hypothetical retrospection (Hansson, 2007a). In this approach, the decision maker concerns different possible outcome scenarios of a decision about acceptable risk and deliberates whether these outcomes scenarios can be deemed acceptable when they materialize (given for example that making a decision about whether or not take a certain risk at a previous point in time was unavoidable). The decision maker can then decide what scenarios should be prevented in deciding about acceptable risk.

In his description of this approach, Hansson (2007a) assumes that the decision maker judges on basis of current moral values, which seems an assumption of moral certainty. However, it would not seem too difficult to adjust the procedure to moral uncertainty. The decision maker might then judge the different hypothetical descriptive outcomes scenarios from the viewpoint of different moral theories or future axiologies (Danaher, 2021), and try to avoid scenarios that are unacceptable on any credible moral theory or normative framework of acceptable risk.[9]

7. Conclusions

Although there is a considerable literature on acceptable risks and various normative frameworks for judging the acceptability of risk impositions, almost none of these approaches have considered moral uncertainty. In this article, I have explored how we can account for moral uncertainty in decisions about acceptable risk. I have rejected MEC (Maximizing Expected Choiceworthiness) as an appropriate approach, mainly because of problems of intertheoretical incomparability and the lack of justifiability of the resulting risk impositions. Instead, I have argued that we should take a more deliberative and coherentist approach to acceptable risk under moral uncertainty. Such an approach would consider moral uncertainty as one of the concerns that needs to be addressed alongside descriptive uncertainty, and other concerns traditionally articulated in risk ethics, like consent, harm, benefits, autonomy, justice and the like.

Although I have focused on decisions about acceptable risk, some of the findings extend beyond that context. Decision-theoretical approaches to moral uncertainty not only face problems of intertheoretical incomparability and fetishist motivation as has been recognized in the literature before, they are also based on problematic assumptions about moral uncertainty and how to deal with it. They ignore moral uncertainty that cannot be expressed in terms of credencs and assume that moral theories are always competing rather than complementary. Moreover, they wrongly assume that all moral theories deal with descriptive and moral uncertainty in a similar way. Although I have not sketched a full-blown alternative approach for decision-making under moral uncertainty, it is clear that there are alternative, more promising approaches than the current decision-theoretical approach on which MEC is based.

Notes

1. An exception is Hayenhjelm (2018).
2. There are also other definitions of risk (Hansson, 2023). In technological contexts, risk is often defined as the probability of an unwanted event times the negative effects of that unwanted event (expected negative value). In decision theory, often a distinction is made between decisions under risk (known probabilities) and decisions under (descriptive) uncertainty (unknown probabilities). Here I do not follow that terminology. In the definition used here, risks are always undesirable but they may nevertheless be acceptable because the activities that introduce risks may also have benefits. As we will see below, questions about what risks are acceptable may also be formulated in terms of the (lack of) acceptability of risking, i.e., taking risks (e.g. mountain climbing) or imposing risks on others (e.g. driving a car in a public road where you may harm others).
3. The list is not meant to be exhaustive. I also kept the formulations short, which might mean that they do not do full justice to the underlying accounts or theories. It is also not always clear whether the mentioned authors intend these principles as the ultimate principles for judging the acceptability of risks and risking or as a main consideration (leaving room for other considerations). My main aim with the list is to show that there are competing ethical frameworks for judging (or deciding about) the acceptability of risks and risking, and that it would not seem far-fetched to assume that decision-makers may well be morally uncertain about which of these is right.
4. Note that this principle is more stringent than the previous one as it also requires considering other activities available for achieving the same ends, and selecting the one with the highest expected value.
5. Although most authors consider risks acceptable if they meet this informed consent principle, several consider the principle too strict, particularly for collective risks because it will lead to a "society of stalemates" (Hansson, 2003, 300).
6. It might be objected that trade-offs as such do not block (weak) justifiability. That is true. It would, for example seem a good (weak) justification to say that there are reasons for and against option A and that - on the balance of reasons – it is still the preferable or an acceptable option. The problem is that credence-weighted trade-offs combine balancing reasons with credences in moral theories, so that it becomes unclear how reasons themselves are weighted, which make justification in terms of underlying moral reasons impossible.
7. This is, for example, suggested by the individuation principle for moral theories proposed in Gustafsson and Torpman (2014) that I discussed above.

8. The interpretation of credences in terms of probabilities, as is done in MEC, is also problematic for more formal reasons, as it debatable whether credences can be meaningfully interpreted as probabilities Côté, N. (2023). There Is No Such Thing as Expected Moral Choice-Worthiness. *Canadian Journal of Philosophy*, *53*(1), 1–20. https://doi.org/10.1017/can.2023.18.
9. Of course, like Hansson (2007a) recognizes, the result may be that there is no or more than one acceptable scenario. If there is no scenario the decision-maker may conclude that the risk should not be taken or when, a decision is unavoidable, they might look for scenarios that or the least unacceptable. I also do not want to deny that at this stage, decision-theoretic approaches might help the decision maker. My point is decisions about acceptable risk under moral uncertainty should not be formulated in the decision-theoretical way assumed by MEC upfront.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This publication is part of the research programme Ethics of Socially Disruptive Technologies, which is funded through the Gravitation programme of the Dutch Ministry of Education, Culture, and Science and the Netherlands Organization for Scientific Research (NWO grant number 024.004.031).

ORCID

Ibo van de Poel http://orcid.org/0000-0002-9553-5651

References

Asveld, L., and S. Roeser, eds. 2009. *The Ethics of Technological Risk*. London:Earthscan.
Beauchamp, T. L., and J. F. Childress. 2013. *Principles of Biomedical Ethics*. 7th ed. Oxford:Oxford University Press.
Côté, N. 2023. "There Is No Such Thing as Expected Moral Choice-Worthiness." *Canadian Journal of Philosophy* 53 (1): 1–20. https://doi.org/10.1017/can.2023.18.
Danaher, J. 2021. "Axiological Futurism: The Systematic Study of the Future of Values." *Futures* 132: 102780. https://doi.org/10.1016/j.futures.2021.102780.
Daniels, N. 1996. *Justice and Justification. Reflective Equilibrium in Theory and Practice*. Cambridge:Cambridge University Press.
Ferretti, M. P. 2016. "Risk Imposition and Freedom." *Politics, Philosophy & Economics* 15 (3): 261–279. https://doi.org/10.1177/1470594X15605437.
Finkelstein, C. 2003. "Is Risk a Harm?" *University of Pennsylvania Law Review* 151 (3): 963–1001. https://doi.org/10.2307/3312883.
Fischhoff, B., S. Lichtenstein, and P. Slovic. 1981. *Acceptable Risk*. Cambridge:Cambridge University Press.
Glickman, T. S., and M. Gough, eds. 1990. *Readings in Risk*. Resources for the Future. Washington.
Gracely, E. J. 1996. "On the Noncomparability of Judgments Made by Different Ethical Theories." *Metaphilosophy* 27 (3): 327–332. https://doi.org/10.1111/j.1467-9973.1996.tb00212.x.
Gustafsson, J. E. 2022. "Second Thoughts About My Favourite Theory." *Pacific Philosophical Quarterly* 103 (3): 448–470. https://doi.org/10.1111/papq.12408.
Gustafsson, J. E., and O. Torpman. 2014. "In Defence of My Favourite Theory." *Pacific Philosophical Quarterly* 95 (2): 159–174. https://doi.org/10.1111/papq.12022.
Haasnoot, M., J. H. Kwakkel, W. E. Walker, and J. ter Maat. 2013. "Dynamic Adaptive Policy Pathways: A Method for Crafting Robust Decisions for a Deeply Uncertain World." *Global Environmental Change* 23 (2): 485–498. https://doi.org/10.1016/j.gloenvcha.2012.12.006.
Hansson, S. O. 2003. "Ethical Criteria of Risk Acceptance." *Erkenntnis* 59 (3): 291–309. https://doi.org/10.1023/A:1026005915919.
Hansson, S. O. 2007a. "Hypothetical Retrospection." *Ethical Theory and Moral Practice* 10 (2): 145–157. https://doi.org/10.1007/s10677-006-9045-3.
Hansson, S. O. 2007b. "Philosophical Problems in Cost-Benefit Analysis." *Economics and Philosophy* 23: 163–183.
Hansson, S. O. 2023. "Risk." In *The Stanford Encyclopedia of Philosophy*, edited by E. N. Zalta. Summer 2023 Edition.

Harman, E. 2015. "The Irrelevance of Moral Uncertainty." In *Oxford Studies in Metaethics*, edited by R. Shafer-Landau, 53–79. Vol. 10. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:Oso/9780198738695.003.0003.
Hayenhjelm, M. 2018. "Risk Impositions, Genuine Losses, and Reparability as a Moral Constraint." *Ethical Perspectives*, 25 (3): 419–446. https://doi.org/10.2143/EP.25.3.3285424.
Hayenhjelm, M., and J. Wolff. 2012. "The Moral Problem of Risk Impositions: A Survey of the Literature." *European Journal of Philosophy* 20 (S1): E26–E51. https://doi.org/10.1111/j.1468-0378.2011.00482.x.
Hedden, B. 2015. "Does MITE Make Right?: On Decision-Making under Normative Uncertainty." In *Oxford Studies in Metaethics*, edited by R. Shafer-Landau, 102–128. Vol. 11. Oxford:Oxford University Press. https://doi.org/10.1093/acprof:Oso/9780198784647.003.0005.
Hicks, A. 2022. "Non-Ideal Prescriptions for the Morally Uncertain." *Philosophical Studies* 179 (4): 1039–1064. https://doi.org/10.1007/s11098-021-01686-1.
Kernohan, A. 2021. "Descriptive Uncertainty and Maximizing Expected Choice-Worthiness." *Ethical Theory and Moral Practice* 24 (1): 197–211. https://doi.org/10.1007/s10677-020-10139-3.
Klinke, A., and O. Renn. 2012. "Adaptive and Integrative Governance on Risk and Uncertainty." *Journal of Risk Research* 15 (3): 273–292. https://doi.org/10.1080/13669877.2011.636838.
Lazar, S., and P. A. Graham. 2021. "Deontological Decision Theory and Lesser-Evil Options." *Synthese* 198 (7): 6889–6916. https://doi.org/10.1007/s11229-019-02496-z.
Lockhart, T. 2000. *Moral Uncertainty and Its Consequences*. Oxford:Oxford University Press.
MacAskill, W., K. Bykvist, and T. Ord. 2020. *Moral Uncertainty*. Oxford:Oxford University Press.
Maheshwari, K., and S. Nyholm. 2022. "Dominating Risk Impositions." *The Journal of Ethics* 26 (4): 613–637. https://doi.org/10.1007/s10892-022-09407-4.
Martin, M. W., and R. Schinzinger. 2005. *Ethics in Engineering*. 4th ed. New York:McGraw-Hill.
Nissan-Rozen, I. 2015. "Against Moral Hedging." *Economics and Philosophy* 31 (3): 349–369. https://doi.org/10.1017/S0266267115000206.
Oberdiek, J. 2017. *Imposing Risk: A Normative Framework*. Oxford:Oxford University Press.
Oddie, G., and P. Milne. 1991. "Act and Value: Expectation and the Representability of Moral Theories." *Theoria* 57 (1–2): 42–76. https://doi.org/10.1111/j.1755-2567.1991.tb00540.x.
Otway, H. J., and D. Von Winterfeldt. 1982. "Beyond Acceptable Risk: On the Social Acceptability of Technologies." *Policy Sciences* 14 (3): 247–256. https://doi.org/10.1007/BF00136399.
Placani, A. 2017. "When the Risk of Harm Harms." *Law and Philosophy* 36 (1): 77–100. https://doi.org/10.1007/s10982-016-9277-x.
Rawls, J. [1971] 1999. *A Theory of Justice*. Revised Edition ed. Cambridge, Massachusetts:The Belknap Press of Harvard University Press.
Riedener, S. 2020. "An Axiomatic Approach to Axiological Uncertainty." *Philosophical Studies* 177 (2): 483–504. https://doi.org/10.1007/s11098-018-1191-7.
Roeser, S., R. Hillerbrand, M. Peterson, and P. Sandin. 2012. *Handbook of Risk Theory: Epistemology, Decision Theory, Ethics, and Social Implications of Risk*. New York:Springer.
Rosenthal, C. 2021. "What Decision Theory Can't Tell us about Moral Uncertainty." *Philosophical Studies* 178 (10): 3085–3105. https://doi.org/10.1007/s11098-020-01571-3.
Ross, J. 2006. "Rejecting Ethical Deflationism." *Ethics* 116 (4): 742–768. https://doi.org/10.1086/505234.
Sepielli, A. 2009. "What to Do When You Don't Know What to Do." In *Oxford Studies in Metaethics*, edited by R. Shafer-Landau, 5–28. Vol. 4. Oxford:Oxford University Press. https://doi.org/10.1093/oso/9780199566303.003.0002.
Sepielli, A. 2016. "Moral Uncertainty and Fetishistic Motivation." *Philosophical Studies* 173 (11): 2951–2968. https://doi.org/10.1007/s11098-016-0645-z.
Shrader-Frechette, K. S. 1991. *Risk and Rationality. Philosophical Foundations for Populist Reform*. Berkeley:University of California Press.
Slovic, P. 2000. *The Perception of Risk*. London:Earthscan.
Smith, M. 1995. *The Moral Problem*. Oxford:Blackwell.
Sunstein, C. R. 2005. "Cost-Benefit Analysis and the Environment." *Ethics* 115 (2): 351–385. https://doi.org/10.1086/426308.
Tarsney, C. 2019. "Normative Uncertainty and Social Choice." *Mind* 128 (512): 1285–1308. https://doi.org/10.1093/mind/fzy051.
Tenenbaum, S. 2017. "Action, Deontology, and Risk: Against the Multiplicative Model." *Ethics* 127 (3): 674–707. https://doi.org/10.1086/690072.
Von Neumann, J., and O. Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton:Princeton University Press.
Weatherson, B. 2014. "Running Risks Morally." *Philosophical Studies* 167 (1): 141–163. https://doi.org/10.1007/s11098-013-0227-2.