CHAPTER 30

Evaluative Uncertainty

JOE HORTON AND JACOB ROSS

1. Introduction

It is not unusual to be uncertain about what is morally valuable, or about what you have moral reason to do. This can happen when you are uncertain about the physical world. For example, you might be uncertain about whether you have more reason to go to the beach or to the library because you are uncertain about whether it is going to rain, or you might be uncertain about whether you have more reason to order the salad or the burger because you are uncertain about what these dishes contain.

But uncertainty about reasons for action need not derive from uncertainty about the physical world. It can instead be fundamental. For example, you might be uncertain about whether you have more reason to go to the beach or to the library because you are uncertain about whether pleasure or knowledge is more valuable, or you might be uncertain about whether you have more reason to order the salad or the burger because you are uncertain about whether nonhuman animals have moral status.

We can think of this fundamental uncertainty as concerning rival *evaluative theories*—theories about values or objective reasons for action. We refer to uncertainty over evaluative theories as *evaluative uncertainty*. An evaluative theory might be comprehensive, covering all values or reasons, or it might be restricted to values or reasons of a particular kind, such as moral or prudential. These values or reasons might be relativised to agents and times. Since evaluative theories include theories of moral reasons, evaluative uncertainty includes moral uncertainty.

In this chapter, we focus on two questions about evaluative uncertainty:

1. How, if at all, does evaluative uncertainty affect the *rationality* of our choices?
2. How, if at all, does evaluative uncertainty affect the *morality* of our choices?

We address the rational question in Sections 2 and 3. We argue that evaluative uncertainty is relevant to the rationality of choice, and we consider several views about how. We then do the same in Sections 4 and 5, but with respect to the moral question. In each case, we critically examine the main views that have been offered in the literature, and we propose an alternative approach.

Most of the views that we consider, including those we favour, face two significant challenges, known as *the problem of intertheoretic value comparison* and *the problem of fanaticism*. We address the former in Section 6 and the latter in Section 7.

It is worth noting that the views in the literature tend to be presented as views about what you rationally or morally ought to do under *normative uncertainty*, where this includes not only evaluative uncertainty but also uncertainty about what you rationally or morally ought to do. But this raises a problem. For, as others have pointed out, it seems incoherent for you to be genuinely uncertain about what you ought to do while also believing that, given this

uncertainty, you ought to act in some particular way.[1] If we instead restrict ourselves to asking about what you rationally or morally ought to do under *evaluative uncertainty*, then we avoid this issue.

For interested readers, we explore the themes of this chapter further, and in distinct ways, in our article 'Evaluative Uncertainty and Permissible Preference'.[2]

2. The Relevance of Evaluative Uncertainty to Rational Choice

It is sometimes argued, or assumed, that evaluative uncertainty is irrelevant to the rationality of choice.[3] There are at least two versions of this view. According to the first version, evaluative uncertainty is irrelevant because, if you are fully rational, then you will never face such uncertainty—you will always be certain of the true evaluative theory. According to the second version, evaluative uncertainty is irrelevant (regardless of whether it is ever rationally permissible) because evaluative belief in general—even full evaluative belief—is irrelevant to the rationality of choice. In this section, we argue against each of these versions in turn.

2.1. The Rationality of Evaluative Uncertainty

There are simple arguments that strongly suggest that, at least for beings like us, evaluative uncertainty is not only rationally permissible but also rationally required. To give such an argument, we need two goods that allow for trade-offs, in the sense that a loss in one can be outweighed by a sufficient gain in the other, and vice versa. In our example, we assume that pleasure and achievement are two such goods.

Let L_0 be an ordinary life containing an ordinary level of pleasure and an ordinary level of achievements. Now, compare L_0 with each life in this sequence:

L_1 is a life exactly like L_0, but with moderately fewer achievements and minutely more pleasure.

L_2 is a life exactly like L_1, but with minutely more pleasure.

L_3 is a life exactly like L_2, but with minutely more pleasure.

. . .

$L_{1,000,000}$ is a life exactly like $L_{999,999}$, but with minutely more pleasure. It contains vastly more pleasure than L_1.

Corresponding to this sequence of lives, there will be a sequence of evaluative theories, T_1 through $T_{1,000,000}$, such that the *n*th theory in this sequence states that L_n is the first life in the sequence of lives that is better, or more choiceworthy, than L_0. If, as we are assuming, pleasure and achievements admit of trade-offs, then one of the theories in this sequence will be true. But clearly none of us is in a position to know which of these theories is true. If we did believe the true theory, that would just be a lucky guess. Since we are not in a position to know which of these theories is true, we should not be certain of any of them, and we should instead divide our credence among several of them. So, it seems that, for beings like us, evaluative uncertainty is rationally required.

It might be objected that it could be vague which life is the first in the sequence that is more choiceworthy than L_0.[4] But this appeal to vagueness would not undermine our argument. If epistemicism is the correct account of vagueness, then one of the corresponding evaluative theories will still be true, although it will be unknowable which. If some other account of

vagueness is correct, then it may be that none of these theories is determinately true. But we could construct a sequence of theories regarding which life is the first in the sequence that is *fully determinately* more choiceworthy than L_o, in the sense that there is no indeterminacy at any level concerning whether this life is more choiceworthy. And we would then be rationally required to be uncertain about which of these theories is true.[5]

2.2. The Relevance of Evaluative Beliefs to Rational Choice

There is a natural way to motivate the view that evaluative belief in general—even full evaluative belief—is irrelevant to the rationality of choice. This motivation comes from cases in which you have crazy evaluative beliefs. Suppose you have the groundless belief that you have most reason to grind your face into broken glass.[6] Presumably, this belief does not make doing so rational. And if even your *full* evaluative beliefs are not relevant to what it is rational for you to do, then it seems your *partial* evaluative beliefs, or credences, cannot be relevant either.

These cases illustrate that we must reject principles like

> **Narrow-Scope Enkrasia**: For any two options A and B, if you believe that you have more reason to choose A than to choose B, then it is irrational for you to choose B over A.

However, such cases do nothing to impugn 'wide-scope' analogues, such as

> **Wide-Scope Enkrasia**: For any two options A and B, it is irrational to choose B over A while believing that you have more reason to choose A than to choose B.[7]

If you believe that you have most reason to grind your face into glass, Narrow-Scope Enkrasia has the implausible implication that you are rationally required to do so. Wide-Scope Enkrasia avoids this implication. It implies only that you cannot rationally refrain from grinding your face into glass *while retaining this belief*. It is thus consistent with the view that you should revise this belief and refrain from grinding your face into glass.

Wide-Scope Enkrasia is widely accepted, for to violate it is to act against your own judgment, and that seems paradigmatically irrational. This suggests that evaluative beliefs are indeed relevant to the rationality of choice—practical rationality seems to require *consistency* between your choices and your evaluative beliefs.

It is natural to expect that this consistency requirement generalises, so as to require consistency between your choices and your evaluative credences. We will now turn to consider some candidates for this generalised requirement. If there is such a requirement, and if, as we argued above, evaluative uncertainty is rationally required, then evaluative uncertainty is relevant to the rationality of choice.

3. Rational Choice Under Evaluative Uncertainty

How does evaluative uncertainty affect the rationality of choice? In this section, we consider three views offered in the literature, and we argue that none of them is entirely satisfactory. We then propose an alternative approach.

There is an important problem with each of the three views from the literature. They each imply that when you have crazy evaluative beliefs, such as the belief that you have most reason to grind your face into glass, you are rationally required to act on them. However, this problem is not decisive. For each of these views could be formulated in broadly the same way as Wide-Scope Enkrasia, and they would then avoid this implication.[8]

We will be considering several examples involving evaluative uncertainty. In order to side-step the problem arising from crazy evaluative beliefs, we will be assuming that, in each example, the stipulated evaluative uncertainty is not crazy but is instead rationally permissible given the evidence available. If you doubt that the evaluative uncertainty in a particular example could ever be rationally permissible, then you could devise a structurally equivalent example involving evaluative uncertainty that seems to you reasonable.

3.1. My Favourite Theory

A natural view, which many people seem to follow in practice, is that you should follow the recommendations of whatever evaluative theory you regard as most plausible. We can state this view more precisely as follows. Let us say that an option is *optimal* if and only if there is no alternative that you have more reason to choose. According to

> **My Favourite Theory (MFT)**: An option is rationally permissible for you if and only if it is optimal according to the evaluative theory, or one of the evaluative theories, in which you have most credence.[9]

However, MFT has some very counterintuitive implications. Suppose that in

> *Near Consensus*: You must decide between a painful operation in the very distant future and a minutely less painful operation today. You have credence 0.1 in a *time-neutral* evaluative theory, according to which the immediate operation is optimal. But the remainder of your credence is divided evenly among ten *future-discounting* evaluative theories, each of which implies that the future operation is optimal.

Since you have credence 0.1 in the time-neutral theory and only credence 0.09 in each of the future-discounting theories, MFT implies that you rationally ought to choose the immediate operation. But that is counterintuitive, for you are very confident that you have more reason to choose the future operation.[10]

3.2. My Favourite Option

Near Consensus suggests that when making decisions under evaluative uncertainty, what matters, fundamentally, is not your credence in different evaluative theories but, rather, your credence in the optimality of your options. So, it might seem that we should adopt

> **My Favourite Option (MFO)**: An option is rationally permissible for you if and only if it is the option, or one of the options, in whose optimality you have most credence.[11]

MFO captures the intuition that, in Near Consensus, you rationally ought to choose the future operation. This is because your credence that the future operation is optimal is 0.9, whereas your credence that the immediate operation is optimal is only 0.1.

However, MFO also has some very counterintuitive implications. Suppose that in

> *Skewed Stakes*: You must choose between a life as a surfer and a life as an academic. You know the surf life would bring slightly more pleasure, whereas the academic life would bring vastly more knowledge. You have credence 0.51 in a *hedonistic* evaluative theory, according to which pleasure is the only thing that matters. But the remainder of your credence is in a *pluralist* evaluative theory, according to which knowledge also matters to a significant degree.

Since you have most credence that the surf life is optimal, MFO implies that you rationally ought to choose the life as a surfer. But that would seem to be a very bad gamble, for you are certain that the surf life is not much better than the academic life, and you think the surf life is reasonably likely to be much worse.[12]

3.3. Maximize Expected Choiceworthiness

Let us say that the *choiceworthiness* of an option is the net strength of all the reasons for and against it. Skewed Stakes suggests that when making decisions under evaluative uncertainty, you need to take into account not only your credence in different evaluative theories but also the degrees of choiceworthiness that these theories ascribe to your options. So, it might seem that we should adopt

> **Maximize Expected Choiceworthiness (MEC)**: An option is rationally permissible for you if and only if it maximizes expected choiceworthiness (where the expected choiceworthiness of an option is calculated by taking each evaluative theory in which you have credence, multiplying your credence in this theory by the choiceworthiness that this theory ascribes to this option, and then summing the results).[13]

MEC has been very influential in the literature. It captures the intuitions that you rationally ought to choose the future operation in Near Consensus and that you rationally ought to choose the life as an academic in Skewed Stakes.

However, MEC also has some very counterintuitive implications. Suppose that in

> *Happy Nihilist*: You are a committed nihilist, certain that no evaluative claims are true and hence that no option has any degree of choiceworthiness. However, despite your nihilism, you have a wonderful life, and you strongly prefer to keep living. You must now decide whether to kill yourself.

Given your strong preference for remaining alive, it seems irrational for you to kill yourself. But MEC cannot account for this. Since you are certain that no option has any degree of choiceworthiness, staying alive cannot have greater expected choiceworthiness than killing yourself.

It might be responded that because nihilist theories fail to ascribe choiceworthiness at all, we should ignore your credence in these theories when applying MEC. However, if we ignore your credence in these theories, a worse problem arises. To see this, consider a version of Happy Nihilist in which, while you are almost certain in nihilism, you also have a smidgeon of credence in an evaluative theory according to which, because all life is sinful, you have most reason to kill yourself. The approach we are considering now not only fails to capture the judgment that it is irrational for you to kill yourself but also implies that killing yourself is rationally required, despite your being almost certain that you have no reason to kill yourself and despite your having a wonderful life that you strongly prefer to keep living.

Happy Nihilist involves an extreme choice. But the problem it raises is very general. For any two options A and B, it seems irrational to choose A over B while preferring B to A, even if you are a committed nihilist. MEC cannot account for this.

It might be irrational to be certain in nihilism. But the same problem arises whenever you have any credence whatsoever in nihilism. This is because, conditional on any nihilist theory, the choiceworthiness of your options is undefined. So, if you have any credence in

any nihilist theory, the expected choiceworthiness of your options is undefined. So, if you have any credence in any nihilist theory, MEC implies that the rational permissibility of your options is undefined.[14]

Furthermore, this problem arises not only when you have any credence in nihilism but also when you have any credence in any evaluative theory that fails to assign determinate degrees of choiceworthiness to your options.[15] This is especially troubling if, as seems very plausible, everyone should have at least a smidgeon of credence in at least some such indeterminate theories. For in that case, MEC implies that rational permissibility is undefined not just for nihilists but also for every rational agent.

3.4. Maximize Expected Utility

Happy Nihilist suggests that when making decisions under evaluative uncertainty, what it is rational for you to do depends not only on your credences but also, non-derivatively, on your preferences. There is an influential view about rational choice on which your preferences have such central importance. It is sometimes called

> **Maximize Expected Utility (MEU)**: It is irrational to choose an option while having credences and preferences relative to which the option fails to maximize expected utility (where utility is a measure of the satisfaction of your preferences).

It is widely accepted that MEU, or something like it, is the correct view about rational choice under *ordinary* uncertainty—uncertainty about what the world is like. But MEU has not been considered in the literature on evaluative uncertainty. Presumably, it has been assumed that because MEU focuses on your preferences rather than on your evaluative beliefs, it is not sensitive to your credences in evaluative theories. We will now argue that this assumption is a mistake, and that MEU works just as well as a view about rational choice under evaluative uncertainty.

The expected utility of an option is calculated by taking each possible situation that you might be in (each way that everything not under your control might be), multiplying your credence that you are in this situation by the utility of the option given this situation, and then summing the results. More formally,

$$\mathrm{EU}(\mathrm{O}) = \sum_{\mathrm{S}} \mathrm{C}(\mathrm{S}) \times \mathrm{U}(\mathrm{O}\,\&\,\mathrm{S})$$

where EU(O) is the expected utility of option O, S ranges over the situations that you might be in, C(S) is your credence that you are in S, and U(O&S) is the utility of O given S.[16]

Since MEU is typically discussed in the context of ordinary uncertainty, the various situations that you might be in are typically specified in purely descriptive terms (e.g., 'it will rain' or 'it will be sunny'). But just as you have no control over certain aspects of the physical world, you likewise have no control over which evaluative theory is true. So, in cases involving both ordinary and evaluative uncertainty, the situations that you could be in must be specified in both descriptive and evaluative terms ('it will rain and only pleasure matters' or 'it will be sunny and both pleasure and knowledge matter'). Since the cases on which we are focusing involve only evaluative uncertainty and not ordinary uncertainty, the relevant situations can be specified in purely evaluative terms ('only pleasure matters' or 'pleasure and knowledge both matter').

If a given situation is defined partly or wholly in evaluative terms, then it seems there will be rational constraints on what your preferences can be conditional on this situation. As an illustration, let S be a situation in which going to the beach would produce more pleasure than going to the library, and in which only pleasure matters. It follows from how S is defined that if S obtains, then you have more reason to go to the beach than to go to the library. So, it follows that if S obtains, then your preferences are fitting only if you prefer going to the beach. Thus, it seems that if you are rational, you will prefer going to the beach *conditional on S*. If we generalise and precisify this idea, we arrive at

> **The Rational Conditional Preference Principle (RCPP)**: For any maximally specific situation S, your preferences conditional on S are rationally permissible if and only if it is consistent with S that these preferences are fitting, in the sense of being adequately responsive to any values or objective reasons entailed by S.

If we accept RCPP, it follows that MEU is indeed sensitive to your credences in evaluative theories, in a way that makes it plausible as a view about rational choice under evaluative uncertainty. As an illustration, return to Near Consensus. RCPP implies that, if you are fully rational, you will *weakly* prefer the immediate operation conditional on the time-neutral theory, and you will *more strongly* prefer the future operation conditional on each of the ten future-discounting theories. Since the vast majority of your credence is divided among the ten future-discounting theories, it follows that, if you are fully rational, only the future operation will maximize expected utility for you. So, MEU captures the intuition that you should choose the future operation.

By similar reasoning, MEU agrees with our intuition about Skewed Stakes. It also agrees with our intuition about Happy Nihilist, and it straightforwardly avoids the implication that rational permissibility is undefined for agents who have credence in indeterminate theories. It is able to do all this because it focuses on your conditional preferences rather than directly on the choiceworthiness of your options, and you can rationally have preferences over options even conditional on indeterminate theories.

It is interesting to note that if we accept MEU, we must reject two views that have been proposed in the literature regarding the practical significance of nihilist theories. The first view holds that credence in nihilist theories renders rational permissibility undefined.[17] MEU differs with this view in the way just explained. The second view holds that, at least in many decision situations, it is rational to *ignore* your credence in nihilist theories because taking this credence into account makes no difference to what it is most rational for you to do.[18] MEU differs with this view because it holds that what it is rational for you to do depends on your conditional preferences, and you can rationally have preferences over options even conditional on nihilism.

Notice that MEU is wide-scope in form, like Wide-Scope Enkrasia. It therefore avoids the implication that it is rational for you to act on crazy beliefs, such as the belief that you have most reason to grind your face into glass. MFT, MFO, and MEC could also be given wide-scope form, so as to likewise avoid this implication.[19] However, reformulating these views in this way would not help them avoid the other problems we have presented.

In summary, MEU is a well-established theory of rational choice under ordinary uncertainty that can be applied equally well to evaluative uncertainty, and it avoids many of the problems facing other theories of rational choice under evaluative uncertainty.[20] In the next two sections, we argue that it can also be developed into a plausible view about moral choice under evaluative uncertainty.

4. The Relevance of Evaluative Uncertainty to Moral Choice

It is sometimes argued, or assumed, that evaluative uncertainty is irrelevant to the morality of choice.[21] There are again at least two versions of this view. According to the first version, evaluative uncertainty is irrelevant because if you are fully moral, you will never face such uncertainty—you will always be certain of the true evaluative theory. According to the second version, evaluative uncertainty is irrelevant (regardless of whether it is ever morally permissible) because evaluative belief in general—even full evaluative belief—is irrelevant to the morality of choice. In this section, we argue against each of these versions in turn.

4.1. The Morality of Evaluative Uncertainty

We have argued that evaluative uncertainty is sometimes rationally required. It might be held that evaluative uncertainty is morally impermissible even when it is rationally required. But that seems implausible, for it seems implausible that morality could require us to have irrational beliefs. So, if we accept that evaluative uncertainty is sometimes rationally required, then we should also accept that evaluative uncertainty is sometimes morally permissible.

4.2. The Relevance of Moral Beliefs to Moral Choice

There is a natural way to motivate the view that evaluative belief in general—even full evaluative belief—is irrelevant to the morality of choice. This motivation comes from cases in which you have crazy moral beliefs.[22] Suppose you have the groundless belief that you have most moral reason to torture kittens. Presumably, this belief does not make doing so moral. And if even your full evaluative beliefs are not relevant to what it is moral for you to do, then it seems your evaluative credences cannot be relevant either.

This position is consistent with Wide-Scope Enkrasia, for the latter is a rational requirement rather than a moral requirement. However, there may be a moral analogue of Wide-Scope Enkrasia with which this position conflicts—or at least is in tension. Just as it seems *rationally* criticisable to choose an option A over an option B while believing that you have more reason to choose B, so it seems *morally* criticisable to choose an option A over an option B while believing that only B is morally permissible. For choosing in this way seems unconscientious.

This moral analogue of Wide-Scope Enkrasia avoids the implausible implication that if you have the crazy belief that you have most moral reason to torture kittens, then you morally ought to do so. It implies only that you cannot morally refrain from torturing kittens *while retaining this crazy belief*. It is thus consistent with the view that you should revise this belief and refrain from torturing kittens.

If there is a moral analogue of Wide-Scope Enkrasia, then again it is natural to expect that this requirement generalises from your beliefs to your credences. And there would then be good grounds for holding that evaluative uncertainty is relevant to moral choice. We will now turn to consider some candidates for this generalised requirement.

5. Moral Choice Under Evaluative Uncertainty

How does evaluative uncertainty affect the morality of choice? In this section, we consider two views offered in the literature, and we argue that neither of them is entirely satisfactory. We then propose an alternative approach, which builds on MEU.

The two views from the literature each imply that when you have crazy moral beliefs, you are morally required to act on them.[23] But, again, this problem is not decisive. Like the views that we considered above, these views could be given wide-scope form, so as to avoid this implication.[24]

5.1. Minimize Expected Wrongness

There are moral analogues of My Favourite Theory and My Favourite Option, but these views face problems analogous to the rational versions.[25] So, we can skip over these views, and we can instead begin with

> **Minimize Expected Wrongness (MEW)**: An option is morally permissible for you if and only if it minimizes expected objective moral wrongness.[26]

This is a natural view, and it has been influential in the literature. But it has three significant flaws.

First, because MEW ignores your credences about nonmoral reasons, including prudential reasons, it is counterintuitively demanding. Suppose that in

> *Donation*: You must decide whether to donate all your wealth to charity. You have credence 0.01 in a very demanding moral theory, according to which not donating all your wealth would be objectively morally wrong. But the remainder of your credence is in a more permissive moral theory, according to which both donating and not donating are morally permissible.

Since you are certain that donating is objectively morally permissible, and not donating might be objectively morally wrong, donating minimizes expected objective moral wrongness. So, MEW implies that you are morally required to make this huge donation, even though you are almost certain that doing so is not objectively morally required.[27]

Second, because MEW ignores your credences about reasons to perform supererogatory acts, it is counterintuitively morally cautious. Suppose that in

> *Surprise Party*: You are considering organising a surprise birthday party for a friend. To keep the party a secret, you would need to tell a few white lies. You have credence 0.99 in a common-sense moral theory, according to which telling these lies is an objectively morally permissible way of carrying out this supererogatory deed. But the remainder of your credence is in a stricter moral theory, according to which telling these lies is very slightly objectively morally wrong.

Since you are certain that not organising the party is objectively morally permissible, and since telling these lies might be very slightly objectively morally wrong, not organising the party minimizes expected objective moral wrongness. So, MEW implies that you are morally required not to organise the party, even though you are almost certain that doing so is both objectively morally permissible and supererogatory.

Third, MEW suffers from a version of the problem that we saw for MEC. If you have any credence in any indeterminate evaluative theory, the expected objective moral wrongness of your options is undefined, and MEW then implies that moral permissibility is undefined. This is especially troubling if, as seems very plausible, everyone should have at least a smidgeon of credence in at least some indeterminate theories.

5.2. Maximize Expected Moral Value

Another natural view, which has also been influential in the literature, is

> **Maximize Expected Moral Value (MEMV)**: An option is morally permissible for you if and only if it maximizes expected moral value (where the moral value of an option is the net strength of all the moral reasons for and against it).[28]

MEMV disagrees with MEW in cases like Surprise Party. But its implication in Surprise Party is no less counterintuitive. Since you are almost certain that you have most moral reason to organise the party, and since you are certain that you do not have strong moral reason to refrain from doing so, organising the party uniquely maximizes expected moral value. So, MEMV implies that you are *morally required* to organise the party, even though you are certain that, objectively speaking, it is morally permissible not to.

MEMV also shares the other two defects of MEW: It is counterintuitively demanding in cases like Donation, and it implies that moral permissibility is undefined for anyone with any credence in any indeterminate evaluative theory.

5.3. Morally Permissible Preference and Faultless Choice

We will now propose an alternative approach, which builds on our discussion of MEU.

We previously argued that for your preferences conditional on a maximally specific situation to be *rationally* permissible, they must be *fitting* conditional on this situation. It likewise seems plausible that for your preferences conditional on a maximally specific situation to be *morally* permissible, they must be *morally fitting* conditional on this situation. For having conditional preferences that are not morally fitting seems unconscientious. Specifically, it seems plausible to adopt

> **The Moral Conditional Preference Principle (MCPP)**: For any maximally specific situation S, your preferences conditional on S are morally permissible if and only if it is consistent with S that these preferences are morally fitting, in the sense of being adequately responsive to any decisive moral reasons entailed by S.

If we accept MEU as the correct account of rational choice, and we accept both RCPP and MCPP, then it seems that the most rational and morally conscientious way to engage in decision-making would be as follows. First, bring your credences in line with the available evidence. Second, bring your preferences in line with the conditional preference principles. Third, choose an option that maximizes expected utility relative to your revised credences and preferences. It thus seems plausible to adopt a view that we can call

> **Faultless Choice (FC)**: You choose in a way that is both fully rational and morally conscientious if and only if your credences are rational, your preferences satisfy the conditional preference principles (together with the standard coherence requirements), and you choose an option because it maximizes expected utility relative to your credences and preferences.

To see the appeal of FC, consider another example. Suppose that in

> *Rescue*: You must rescue either Lifeboat A or Lifeboat B. You know that A and B each hold ten children, and B also holds a mouse. You have credence 0.49 in a *taurekian* moral theory, according to which there is a very strong moral reason to give everyone an equal chance of survival and hence to decide between A and B by tossing a coin. But the remainder of your credence is in a *utilitarian* moral theory, according to which, because B contains the mouse, there is very slightly more moral reason to rescue B than to toss a coin.[29]

MCPP implies that, conditional on the utilitarian theory, you are morally required to *slightly* prefer saving B, and conditional on the taurekian theory, you are morally required to *strongly* prefer tossing a coin. If you have these preferences, then tossing a coin uniquely maximizes expected utility, given your credences. FC thus implies that tossing a coin is the only option that could be chosen in a manner that is both fully rational and morally conscientious.

This approach to moral choice under evaluative uncertainty avoids the problems we have seen for MEW and MEMV. To see this, return to Donation. MCPP implies that, conditional on the very demanding moral theory, you are morally required to strongly prefer donating. However, conditional on the more permissive moral theory, a range of preferences are morally permissible, including both a strong preference *for* donating and a strong preference *against* donating. If you strongly prefer to donate conditional on the very demanding theory, and you have at most a very weak preference against donating conditional on the more permissive theory, then donating will maximize expected utility, given your credences. If instead you have a stronger preference against donating conditional on the more permissive theory, then, even with your strong preference for donating conditional on the very demanding theory, *not* donating will maximize expected utility, given your credences. FC thus implies that donating and not donating are both options that could be chosen in a manner that is fully rational and morally conscientious.

Return next to Surprise Party. MCPP implies that, conditional on the strict moral theory, you are morally required to at least weakly prefer not organising the party. However, conditional on the common-sense moral theory, a range of preferences are morally permissible, including both a strong preference *for* organising the party and a strong preference *against* organising the party. If you weakly prefer not organising the party conditional on the strict theory, and your preference for organising the party conditional on the common-sense theory is zero or sufficiently close to zero, then not organising the party will maximize expected utility, given your credences. If instead you have a stronger preference for organising the party conditional on the common-sense theory, then, even with your weak preference against organising the party conditional on the strict theory, organising the party will maximize expected utility, given your credences. FC thus implies that organising the party and not organising the party are both options that could be chosen in a manner that is fully rational and morally conscientious.

Furthermore, whereas MEW and MEMV are unable to provide guidance in cases in which you have credence in indeterminate evaluative theories, FC does not have this drawback. This is because FC takes into account your preferences conditional on indeterminate evaluative theories. And since FC holds that fully rational and morally conscientious choice requires rational credence, it also avoids the implication that when you have crazy beliefs, you should act on them.

There is a further reason to like this kind of approach. It seems plausible that the requirements of rationality and of morality can always be jointly satisfied, so that you are never forced to choose between being rational and being moral. To avoid this kind of

conflict, our account of moral choice must be closely aligned with our account of rational choice. Thus, if we accept MEU as the correct account of rational choice, then we cannot accept a very different account of moral choice, such as MEW or MEMV, but must instead accept something more closely aligned to MEU. FC shows how this can be done.

As we mentioned in the introduction, most of the views we have considered face two further challenges, known as the *problem of intertheoretic value comparison* and the *problem of fanaticism*. These challenges apply to any view about choice under evaluative uncertainty that involves maximizing or minimizing some expectation, such as expected value, expected wrongness, or expected utility. We now turn to consider these challenges.

6. The Problem of Intertheoretic Value Comparison

As we will see in this section, there are several versions of the problem of intertheoretic value comparison. We consider them in turn. The first version arises with respect to all evaluative theories, whereas the other two versions arise only with respect to specific kinds of evaluative theory.

6.1. The General Problem

Return to Rescue. Most of the views about choice under evaluative uncertainty that we have considered imply that in determining what you ought to do in Rescue, you must take into account not only your *credences* in the two evaluative theories but also *how much is at stake* according to each theory. In particular, they imply that if much more is at stake according to the taurekian theory, then you ought to follow its recommendations. Thus, most of these views imply that in determining what you ought to do in this situation, you must compare the value differences (or the strengths of objective reasons) ascribed by the two theories. That is, you must make an *intertheoretic value comparison*.

Some philosophers reject the intelligibility of such comparisons. They hold that although it makes perfect sense to compare value differences *within* the value ordering given by a particular evaluative theory, it makes no sense to compare value differences *across* the value orderings given by different evaluative theories.[30] In other words, on this view, it makes sense to say that, according to a given evaluative theory, the stakes are higher in one choice situation than in another, but it makes no sense to say that, in a given choice situation, the stakes are higher according to one evaluative theory than according to another. Proponents of this view hold that there is no common scale on which to compare the value orderings corresponding to different evaluative theories.

6.2. The Counterintuitiveness of Intertheoretic Incomparability

However, scepticism about intertheoretic value comparisons is counterintuitive. To see this clearly, consider another example. Suppose that in

> *Dinner Party*: You must decide whether to serve roast lamb or vegan risotto, and you know that your guests would very slightly prefer the lamb. You have credence 0.51 in an anthropocentric moral theory, according to which, because only human interests matter, it is very slightly morally wrong to serve the risotto. But the remainder of your credence is in a

non-anthropocentric moral theory, according to which, because all sentient beings matter equally, it is very morally wrong to serve the lamb.

It seems very natural to say that, if all sentient beings matter equally, then there is a lot at stake in Dinner Party, whereas if only human interests matter, then there is much less at stake. But to say this is to make an intertheoretic value comparison.

6.3. The Fitting-Attitudes Solution

Since we have the intuition that intertheoretic value comparisons are meaningful, we should try to provide an account of their meaning. To provide such an account, we need just two assumptions. First, we need to assume that *complete* evaluative theories have implications not only about what you have reason to do but also about the preferences that it is fitting for you to have.[31]

Second, we need to assume that mental states are comparable, not only for a particular person at a particular time but also across people and times. For at least some mental states, this interpersonal and intertemporal comparability seems very plausible. For example, it seems very plausible that just as the pain you are now experiencing from smashing your thumb with a hammer can be greater than the pain you are now experiencing from pricking your finger, so the pain *you* are *now* experiencing from smashing your thumb can be greater than the pain *I* experienced *yesterday* from pricking my finger. Similarly, it seems very plausible that just as the guilt you are now experiencing for cheating on your spouse can be greater than the guilt you are now experiencing for littering, so the guilt *you* are experiencing *now* for cheating on your spouse can be greater than the guilt *I* experienced *yesterday* for littering.

If these affective mental states are comparable across people and times, then the same is plausibly true of preferences. For there appears to be a tight connection between fitting preferences and fitting affects. In particular, if outcomes A and B are exhaustive and mutually exclusive, and if you initially regard them as equally probable, then the more strongly it is fitting for you to prefer A to B, the more pleased it is fitting for you to be upon learning that A obtains, and the more disappointed it is fitting for you to be upon learning that B obtains.[32] But if the strengths of fitting preferences are associated with the intensities of fitting affects, and if, as we have argued, the intensities of affects are comparable across people and times, then the strengths of preferences must likewise be comparable across people and times. So, it seems there must be a universal scale on which the strengths of preferences can be measured.

If complete evaluative theories indicate which preferences are fitting, and there is an independent scale on which preferences can be measured, then intertheoretic value comparisons will be perfectly intelligible. For example, consider the claim that, in Dinner Party, there is more at stake according to the non-anthropocentric theory than there is according to the anthropocentric theory. This can be understood to mean that a stronger preference between your options is fitting according to the non-anthropocentric theory than according to the anthropocentric theory.

We have focused on fitting *preferences*. But one could also attempt to solve the problem of intertheoretic value comparison by appealing to other kinds of fitting attitude. For example, if degrees of moral wrongness correspond to fitting degrees of guilt and anger, so that complete moral theories have implications about these fitting attitudes, then these implications will enable us to compare degrees of moral wrongness across moral theories, so long as there is a universal scale on which degrees of guilt and anger can be measured. So, the solution that we have presented here, which appeals to fitting preferences, is just one instance of a broader solution, which we may call *the fitting-attitudes solution*.

6.4. Incomplete Theories

We have argued that evaluative theories are comparable so long as they are complete, in the sense that they indicate not only what you have most reason to choose but also which preferences it is fitting for you to have. It might be objected that not all evaluative theories are complete in this sense.[33] For example, consider utilitarianism. This theory implies that, other things being equal, killing ten people is ten times as bad as killing one, and hence that your preference for killing no one over killing ten people should be ten times as strong as your preference for killing no one over killing one person. However, these implications concern only the *relative* strengths of fitting preferences, not their *absolute* strengths.

This is true. But what it shows is only that utilitarianism, in its standard formulation, is an incomplete theory that can be precisified in infinitely many ways. There is a precisification of utilitarianism according to which it is fitting to have a *strong* preference for avoiding killing one person, and an even *stronger* preference (by a factor of ten) for avoiding killing ten people. And there is another precisification of utilitarianism according to which it is fitting to have a *mild* preference for avoiding killing ten people, and an even *milder* preference (by a factor of ten) for avoiding killing one person. These precisifications agree about the ratios among the strengths of fitting preferences, but they disagree about their absolute strengths, and so they are distinct theories. So, utilitarianism, in its abstract form, is equivalent to the disjunction of its infinitely many precisifications.

Moreover, we regard some of these precisifications as more plausible than others. Most people are very confident that it is fitting to *strongly* prefer to avoid mass murder, and not fitting to only *mildly* prefer to avoid mass murder.[34] So, most people will have greater credence in precisifications of utilitarianism according to which this preference should be strong than in precisifications according to which it should be weak. And, unlike the abstract formulation of utilitarianism, these precisifications specify the value differences between options in noncomparative terms. So, insofar as our credence is spread over these precisifications, intertheoretic value comparisons will be possible.

6.5. Indeterminate Theories

We have just argued that many incomplete evaluative theories can be thought of as disjunctions of complete evaluative theories, with each disjunct specifying how strongly it is fitting to prefer any given option to any other. But some theories might instead be *indeterminate*. That is, a theory might imply, for at least some pair of options, that no particular strength of preference between them is uniquely fitting.

There are several ways in which an evaluative theory might be indeterminate. First, there are what we may call *fuzzy* theories, which imply that your preferences are fitting so long as their strengths lie within a certain range. Second, there are what we may call *purely relational* theories, which imply that your preferences are fitting so long as the right ratios obtain among their strengths. Third, there are what we may call *purely ordinal* theories, which imply that your preferences are fitting so long as you rank your options appropriately, regardless of how strongly you prefer one to another. Fourth, there are *nihilist* theories, according to which no set of preferences is any more fitting than any other.

How are you to decide what to do when you have credence in indeterminate theories? What makes this question particularly challenging is that your credence may be divided among different kinds of evaluative theories. For example, you might have some credence in

a fully determinate theory, some credence in a purely relational theory, and some credence in a purely ordinal theory.

A number of solutions to this problem have been proposed. They can generally be divided into two types.[35] The first type involves *structural depletion*, in the sense that it involves treating more determinate theories as though they were less determinate. For example, we might ignore the cardinal information provided by evaluative theories and treat them all as though they were purely ordinal. We could then aggregate these theories using an ordinal voting rule.[36] The second type of solution involves *structural enrichment*, in the sense that it involves treating less-determinate theories as though they were more determinate. For example, we might transform purely relational theories into fully determinate theories by treating them as having a range equivalent to the greatest range of any fully determinate theory.[37] We could then aggregate these theories by taking an expectation over them.

These approaches are both questionable.[38] If we adopt structural depletion, then we will sometimes ignore the cardinal information provided by fully determinate theories. But this information seems relevant. If we instead adopt structural enrichment, then we must choose one of the infinitely many ways of enriching indeterminate theories. However, there does not seem to be any justification for choosing any one of these ways rather than any other, and so our procedure will inevitably seem unmotivated.

There is another problem that applies to approaches of both kinds. In cases in which our credence is divided among purely ordinal theories, an approach of either kind will amount to an ordinal voting rule. But it follows from Arrow's Theorem that using any ordinal voting rule will either lead to violations of the Independence of Irrelevant Alternatives or else it will have some other defect that is even worse (e.g., treating one evaluative theory as a 'dictator' or ignoring unanimity across evaluative theories).[39]

If we adopt MEU and FC, then there is no need to engage in structural depletion or structural enrichment, or to use any ordinal voting rule. Instead, we just need to take an expectation over the strengths of your preferences conditional on the evaluative theories in which you have credence. If an evaluative theory is indeterminate, it imposes weaker constraints on your preferences than a fully determinate theory. But presumably your preferences conditional on an indeterminate theory can still have *some particular strength or other*. So, even when you have credence in indeterminate evaluative theories, your options can still have well-defined expected utilities.

7. The Problem of Fanaticism

In this section, we address the problem of fanaticism. As with the problem of intertheoretic value comparison, this problem applies to any view about choice under evaluative uncertainty that involves maximizing or minimizing some expectation, such as expected value, expected wrongness, or expected utility.

7.1. The Problem

Suppose that in

> *Lying for Lives*: You must decide whether to tell a lie to save a thousand lives. You have credence 0.99 in a utilitarian evaluative theory, according to which you have strong reason to tell the lie. But the remainder of your credence is in a fanatical evaluative theory, according to which you have overriding reason to refrain from lying, no matter how many lives are at stake.

It is natural to construe the fanatical theory as assigning infinite disvalue to lying. But the utilitarian theory presumably assigns only finite disvalue to failing to save the thousand lives. If refraining from lying is at most finitely bad, and there is some chance, however small, that lying is infinitely bad, then refraining from lying will maximize expected value. So, if you ought to maximize expected value, then you ought to refrain from lying and allow the thousand lives to perish, even though you are almost certain that doing so is objectively very wrong.[40]

Cases like Lying for Lives clearly raise a problem for value-maximizing views, such as MEC and MEMV, and for wrongness-minimizing views, such as MEW. But they also raise a problem for utility-maximizing views, such as MEU and FC, so long as these views are combined with the conditional preference principles.

Suppose next that, in addition to having credence in an evaluative theory according to which lying has infinite *negative* value (or is infinitely *less* valuable than the alternative), you also have credence in an evaluative theory according to which lying has infinite *positive* value (or is infinitely *more* valuable than the alternative). Since there is some chance that lying has infinite negative value and some chance that lying has infinite positive value, the expected value of lying is undefined.[41] So, if rational and moral requirements are to be understood in terms of the maximization of expected value, then your rational and moral requirements will be undefined.

It might be argued that for every agent and every option, it is irrational for the agent not to have at least a smidgeon of credence both in an evaluative theory according to which this option has infinite positive value and in an evaluative theory according to which this option has infinite negative value. If this is right, and rational and moral requirements are to be understood in terms of the maximization of expected value, then rational and moral requirements will be undefined for every rational agent.

7.2. Two Potential Solutions

There are at least two kinds of potential solution to the problem of fanaticism. First, we could appeal to constraints on rational credences. For example, we could argue that, in cases like Lying for Lives, it is irrational to have more than extremely low credence in evaluative theories that ascribe extreme values or extremely strong reasons, and it is irrational to have more than zero or infinitesimal credence in evaluative theories that ascribe infinite values or infinitely strong reasons.[42] Second, we could revise the relevant decision rule. For example, we could maintain that when your credence in a possibility is sufficiently low, you rationally ought to ignore this possibility in your decision-making.[43]

It might seem difficult to motivate either of these potential solutions. However, versions of the problem of fanaticism can also arise in cases of choice under *ordinary* uncertainty, as illustrated by *Pascal's Wager* and the *St. Petersburg Paradox*.[44] So, the problem of fanaticism is really a problem for everyone, regardless of their views about choice under evaluative uncertainty. Thus, there are contexts in which everyone will need to appeal to at least one of these solutions, or to some similar solution.[45]

There may be independent grounds for embracing the solution that appeals to constraints on rational credences. For it has been argued that ascribing infinite values or having unbounded utilities makes you predictably exploitable by means of a 'Dutch Book'.[46] Since vulnerability to a Dutch Book is often regarded as an indicator of incoherence, this suggests that it may be irrational to have more than zero or infinitesimal credence in evaluative theories that ascribe infinite values or infinitely strong reasons.[47]

8. Conclusion

We have argued for two main conclusions. First, evaluative uncertainty affects the rationality of choice, and it is plausible that it affects the morality of choice as well, and so we seem to need a theory of rational and moral choice under evaluative uncertainty. Second, we may not need any *special* theory, because the dominant theory of rational choice under ordinary uncertainty can be extended to rational and moral choice under evaluative uncertainty.

Acknowledgments

We benefited from discussing the ideas in this chapter with many excellent philosophers, but we owe special thanks to Abelard Podgorski and Connie Rosati.

Notes

1. For discussion of this problem, and proposals for addressing it, see Brian Weatherson, 'Running Risks Morally', *Philosophical Studies* 167 (2014): 141–163; Andrew Sepielli, 'What to Do When You Don't Know What to Do When You Don't Know What to Do . . .', *Noûs* 48 (2014): 521–544; Elizabeth Harman, 'The Irrelevance of Moral Uncertainty', *Oxford Studies in Metaethics* 10 (2015): 53–79; Andrew Sepielli, 'How Moral Uncertaintism Can Be Both True and Interesting', *Oxford Studies in Normative Ethics* 7 (2018): 98–116; and William MacAskill, Krister Bykvist, and Toby Ord, *Moral Uncertainty* (Oxford, UK: Oxford University Press, 2020), 30–33.
2. Joe Horton and Jacob Ross, 'Evaluative Uncertainty and Permissible Preference', *The Philosophical Review* 134 (2025): 35–64.
3. This view is defended in Brian Hedden, 'Does MITE Make Right? On Decision Making Under Normative Uncertainty', *Oxford Studies in Metaethics* 11 (2016): 102–128. See also Ralph Wedgwood, 'Moral Disagreement and Inexcusable Irrationality', *American Philosophical Quarterly* 56 (2019): 97–108.
4. For discussion of vagueness, see Timothy Williamson, *Vagueness* (New York: Routledge, 1994).
5. Notice that, as we have defined *fully determinate*, it cannot be indeterminate whether a given life is fully determinately more choiceworthy than L_o. For if this were indeterminate, then there would be a kind of higher-order indeterminacy concerning whether this life is more choiceworthy than L_o. And from this it would follow that it is *not* fully determinate whether this life is more choiceworthy than L_o. And so a contradiction arises if we assume that it is indeterminate whether a given life is fully determinately more choiceworthy than L_o.
6. We owe this example to Larry Temkin.
7. For a discussion of Enkrasia, see John Broome, 'Enkrasia', *Organon* 20 (2013): 425–436. In order to avoid counterexamples involving uncertainty, Enkrasia must be formulated in terms of maximally specific options. Donald Regan presents one such counterexample in *Utilitarianism and Cooperation* (Oxford, UK: Oxford University Press, 1980), 265. For a structurally equivalent example, see Derek Parfit, *On What Matters, Volume 1* (Oxford, UK: Oxford University Press, 2011), 159. For discussion of how to formulate Enkrasia in light of these counterexamples, see Jacob Ross, "Acceptance and Practical Reason" (PhD Dissertation), 176–177.
8. See note 19 for the reformulated views.
9. A version of this view is advocated in Edward J. Gracely, 'On the Noncomparability of Judgments Made by Different Ethical Theories', *Metaphilosophy* 27 (1996): 327–332; and Johan E. Gustafsson and Olle Torpman, 'In Defence of My Favourite Theory', *Pacific Philosophical Quarterly* 95 (2014): 159–174.
10. A similar objection is pressed against MFT in MacAskill, Bykvist, and Ord, *Moral Uncertainty*, 42–43.

11. Ted Lockhart suggests this view, but goes on to reject it, in *Moral Uncertainty and Its Consequences* (Oxford, UK: Oxford University Press, 2000), 26.
12. A similar objection is pressed against MFO in Lockhart, *Moral Uncertainty and Its Consequences*, 28; and MacAskill, Bykvist, and Ord, *Moral Uncertainty*, 44–47.
13. A version of this view is advocated in Jacob Ross, 'Rejecting Ethical Deflationism', *Ethics* 116 (2006): 742–768; Andrew Sepielli, 'What to Do When You Don't Know What to Do', *Oxford Studies in Metaethics* 4 (2009): 5–28; and MacAskill, Bykvist, and Ord, *Moral Uncertainty*.
14. William MacAskill develops this problem in 'The Infectiousness of Nihilism', *Ethics* 123 (2013): 508–520.
15. This includes evaluative theories according to which your options are incomparable or incommensurable. We discuss indeterminate evaluative theories in more detail in Section 6.4.
16. This is the *causal* formula for EU, but everything we say is compatible with the *evidential* formula.
17. See MacAskill, 'The Infectiousness of Nihilism'.
18. See Ross, 'Rejecting Ethical Deflationism'.
19. A wide-scope formulation of MFT would hold that it is irrational to choose an option while having credences relative to which the option is not optimal according to any of the maximally probable evaluative theories. A wide-scope formulation of MFO would hold that it is irrational to choose an option while having credences relative to which there is an alternative that has a greater probability of being optimal. A wide-scope formulation of MEC would hold that it is irrational to choose an option while having credences relative to which the option fails to maximize expected choiceworthiness.
20. MEU has another advantage over MEC. Abelard Podgorski has argued that MEC needs to be modified to avoid counterexamples in which there is a dependence relation between your descriptive and evaluative uncertainty. See his 'Normative Uncertainty and the Dependence Problem', *Mind* 129 (2020): 43–70. MEU is already sensitive to any such dependence relation, so it avoids these counterexamples without any modification. See also Pamela Robinson, 'Is Normative Uncertainty Irrelevant if Your Descriptive Uncertainty Depends on It?', *Pacific Philosophical Quarterly* 103 (2022): 874–899.
21. This view is defended in Harman, 'The Irrelevance of Moral Uncertainty'; and Hedden, 'Does MITE Make Right?'
22. See, for example, Harman, 'The Irrelevance of Moral Uncertainty'; and Julia Staffel, 'Normative Uncertainty and Probabilistic Moral Knowledge', *Synthese* 198 (2021): 6739–6765. For further discussion of such cases, see Nomy Arpaly, 'Moral Worth', *Journal of Philosophy* 99 (2002): 223–245; Elizabeth Harman, 'Does Moral Ignorance Exculpate?', *Ratio* 24 (2011): 443–468; and Nomy Arpaly, 'Huckleberry Finn Revisited: Inverse Akrasia and Moral Ignorance', in Randolph Clarke, Michael McKenna, and Angela M. Smith (eds.), *The Nature of Moral Responsibility: New Essays* (Oxford, UK: Oxford University Press, 2015), 141–156.
23. Harman develops this problem in 'The Irrelevance of Moral Uncertainty'.
24. See note 19 for wide-scope formulations of the earlier views.
25. To generate these problems, we just need to revise cases like Near Consensus and Skewed Stakes such that they involve moral decisions and moral uncertainty.
26. Lockhart advocates a version of this view in *Moral Uncertainty and Its Consequences*. Notice that this view appeals to 'objective' moral wrongness. This is because, in order to avoid incoherence, we must distinguish the deontic operators that appear on either side of the biconditional.
27. A similar objection is pressed against views like MEW in Brian Weatherson, 'Review: Moral Uncertainty and Its Consequences', *Mind* 111 (2002): 693–696; Christian Barry and Patrick Tomlin, 'Moral Uncertainty and Permissibility: Evaluating Option Sets', *Canadian Journal of Philosophy* 46 (2016): 898–923; Hedden, 'Does MITE Make Right?', 115–116; and Chelsea Rosenthal, 'What Decision Theory Can't Tell Us About Moral Uncertainty', *Philosophical Studies* 178 (2021): 3085–3105.
28. Harman takes this to be the dominant view about moral choice under evaluative uncertainty, and argues against it, in 'The Irrelevance of Moral Uncertainty'.

29. This example is inspired by Leora Sung, 'Never Just Save the Few', *Utilitas* 34 (2022): 275–288.
30. This problem was introduced in James L. Hudson, 'Subjectivization in Ethics', *American Philosophical Quarterly* 26 (1989): 221–229. For detailed discussion, see Hedden, 'Does MITE Make Right?', 107–115; Amelia Hicks, 'Moral Uncertainty and Value Comparison', *Oxford Studies in Metaethics* 13 (2018): 161–183; and MacAskill, Bykvist, and Ord, *Moral Uncertainty*, Chapter 5.
31. We here develop a suggestion made in Ross, 'Rejecting Ethical Deflationism', 765. Ross also suggests that we can define a common unit of value across at least some evaluative theories by focusing on "cases in which, for some pair of options, we know that the difference between their values is the same according to both . . . theories". See 'Rejecting Ethical Deflationism', 764–765. Sepielli develops this suggestion in 'What to Do When You Don't Know What to Do'. However, this suggestion is not going to help where evaluative theories have no common ground. Furthermore, it has been argued that this approach yields contradictions. See Hedden, 'Does MITE Make Right?', 111–112; and MacAskill, Bykvist, and Ord, *Moral Uncertainty*, 130–138.
32. See Tim Schroeder, 'Desire', in Edward N. Zalta (ed.), *Stanford Encyclopedia of Philosophy* (Summer 2020); and Maria Alvarez, 'Desires, Dispositions and the Explanation of Action', in Julien A. Deonna and Federico Lauria (eds.), *The Nature of Desire* (Oxford, UK: Oxford University Press, 2017), 119–136.
33. See Hedden, 'Does MITE Make Right?', 107–115.
34. Let P_1 be a possible person who is indifferent across all alternatives. Let P_n be an ordinary person with ordinary preferences. There is a spectrum of possible persons connecting P_1 and P_n such that everyone on this spectrum after P_1 has the same ratios among the strengths of their preferences, but their preferences differ in absolute strength. In the early part of this spectrum, we find individuals who are *almost indifferent* across all alternatives (and who, correspondingly, are never very pleased or displeased to learn anything). These people prefer not to commit mass murder, but their preference is weak. Most people have the intuition that such a preference is not fitting.
35. For very helpful discussions of these two approaches, see MacAskill, Bykvist, and Ord, *Moral Uncertainty*, Chapter 4; and Christian Tarsney, 'Vive la Différence? Structural Diversity as a Challenge for Metanormative Theories', *Ethics* 131 (2021): 151–182. For alternative approaches, see Jennifer Rose Carr, 'Normative Uncertainty Without Theories', *Australasian Journal of Philosophy* 98 (2020): 747–762; and Jennifer Rose Carr, 'The Hard Problem of Intertheoretic Value Comparisons', *Philosophical Studies* 179 (2022), 1401–1427.
36. For discussion of which voting rule would be most appropriate, see William MacAskill, 'Normative Uncertainty as a Voting Problem', *Mind* 125 (2016): 967–1004; Christian Tarsney, 'Normative Uncertainty and Social Choice', *Mind* 128 (2019): 1285–1308; and MacAskill, Bykvist, and Ord, *Moral Uncertainty*, Chapter 4.
37. A similar approach is suggested in Lockhart, *Moral Uncertainty and Its Consequences*, 84.
38. Tarsney presses these problems in 'Vive la Différence?' For a defence of structural enrichment, which aims to justify a particular way of cardinalizing ordinal theories, see MacAskill, Bykvist, and Ord, *Moral Uncertainty*, Chapter 4.
39. Kenneth J. Arrow, 'A Difficulty in the Concept of Social Welfare', *Journal of Political Economy* 58 (1950): 328–346.
40. This problem was first articulated in Ross, 'Rejecting Ethical Deflationism', 765.
41. This point is emphasised in MacAskill, Bykvist, and Ord, *Moral Uncertainty*, 152–153.
42. For defence of this proposal, see Ross, 'Rejecting Ethical Deflationism', 765–767.
43. For defence of this proposal, see Nicholas J. J. Smith, 'Is Evaluative Compositionality a Requirement of Rationality?', *Mind* 123 (2014): 457–502; Lara Buchak, *Risk and Rationality* (Oxford, UK: Oxford University Press, 2014), 73–74; and Bradley Monton, 'How to Avoid Maximising Expected Utility', *Philosophers' Imprint* 19 (2019): 1–25.
44. We here follow MacAskill, Bykvist, and Ord, *Moral Uncertainty*, 153.
45. Of course, both these potential solutions face challenges. This is an area where further work is needed. For relevant discussion, see Alan Hájek, 'Unexpected Expectations', *Mind* 123 (2014):

533–567; Yoaav Isaacs, 'Probabilities Cannot Be Rationally Neglected', *Mind* 125 (2016): 759–762; and Hayden Wilkinson, 'In Defence of Fanaticism', *Ethics* 132 (2022): 445–477.

46. See Vann McGee, 'An Airtight Dutch Book', *Analysis* 59 (1999): 257–265. This argument is controversial because it involves a countably infinite number of bets. For a critique of this kind of argument, see Frank Arntzenius, Adam Elga, and John Hawthorne, 'Bayesianism, Infinite Decisions, and Binding', *Mind* 113 (2004): 251–283. For a defence of this kind of argument, see Jacob Ross, 'Sleeping Beauty, Countable Additivity, and Rational Dilemmas', *The Philosophical Review* 119 (2010): 411–447.
47. This view of Dutch Book arguments traces back to Frank Ramsey. See 'Truth and Probability', in Frank Ramsey, *The Foundations of Mathematics and Other Logical Essays* (London: Routledge, 1931). See also Colin Howson and Peter Urbach, *Scientific Reasoning: The Bayesian Approach* (Chicago: Open Court, 2006); and David Christensen, *Putting Logic in Its Place: Formal Constraints on Rational Belief* (Oxford, UK: Oxford University Press, 2004).

References

Alvarez, Maria. 'Desires, Dispositions and the Explanation of Action'. In Julien A. Deonna and Federico Lauria (eds.), *The Nature of Desire* (Oxford, UK: Oxford University Press, 2017), 119–136.

Arntzenius, Frank, Adam Elga, and John Hawthorne. 'Bayesianism, Infinite Decisions, and Binding'. *Mind* 113 (2004): 251–283.

Arpaly, Nomy. 'Moral Worth'. *Journal of Philosophy* 99 (2002): 223–245.

Arpaly, Nomy. 'Huckleberry Finn Revisited: Inverse Akrasia and Moral Ignorance'. In Randolph Clarke, Michael McKenna, and Angela M. Smith (eds.), *The Nature of Moral Responsibility: New Essays* (Oxford, UK: Oxford University Press, 2015), 141–156.

Arrow, Kenneth J. 'A Difficulty in the Concept of Social Welfare'. *Journal of Political Economy* 58 (1950): 328–346.

Barry, Christian, and Patrick Tomlin. 'Moral Uncertainty and Permissibility: Evaluating Option Sets'. *Canadian Journal of Philosophy* 46 (2016): 898–923.

Broome, John. 'Enkrasia'. *Organon* 20 (2013): 425–436.

Buchak, Lara. *Risk and Rationality* (Oxford, UK: Oxford University Press, 2014).

Carr, Jennifer Rose. 'Normative Uncertainty Without Theories'. *Australasian Journal of Philosophy* 98 (2020): 747–762.

Carr, Jennifer Rose. 'The Hard Problem of Intertheoretic Value Comparisons'. *Philosophical Studies* 179 (2022): 1401–1427.

Christensen, David. *Putting Logic in Its Place: Formal Constraints on Rational Belief* (Oxford, UK: Oxford University Press, 2004).

Gracely, Edward J. 'On the Noncomparability of Judgments Made by Different Ethical Theories'. *Metaphilosophy* 27 (1996): 327–332.

Gustafsson, Johan E., and Olle Torpman. 'In Defence of My Favourite Theory'. *Pacific Philosophical Quarterly* 95 (2014): 159–174.

Hájek, Alan. 'Unexpected Expectations'. *Mind* 123 (2014): 533–567.

Harman, Elizabeth. 'Does Moral Ignorance Exculpate?' *Ratio* 24 (2011): 443–468.

Harman, Elizabeth. 'The Irrelevance of Moral Uncertainty'. *Oxford Studies in Metaethics* 10 (2015): 53–79.

Hedden, Brian. 'Does MITE Make Right? On Decision Making Under Normative Uncertainty'. *Oxford Studies in Metaethics* 11 (2016): 102–128.

Horton, Joe, and Jacob Ross. 'Evaluative Uncertainty and Permissible Preference'. *The Philosophical Review* 134 (2025): 35–64.

Howson, Colin, and Peter Urbach. *Scientific Reasoning: The Bayesian Approach* (Chicago: Open Court, 2006).

Hudson, James L. 'Subjectivization in Ethics'. *American Philosophical Quarterly* 26 (1989): 221–229.

Isaacs, Yoaav. 'Probabilities Cannot Be Rationally Neglected'. *Mind* 125 (2016): 759–762.

Lockhart, Ted. *Moral Uncertainty and Its Consequences* (Oxford, UK: Oxford University Press, 2000).
MacAskill, William. 'The Infectiousness of Nihilism'. *Ethics* 123 (2013): 508–520.
MacAskill, William. 'Normative Uncertainty as a Voting Problem'. *Mind* 125 (2016): 967–1004.
MacAskill, William, Krister Bykvist, and Toby Ord. *Moral Uncertainty* (Oxford, UK: Oxford University Press, 2020).
McGee, Vann. 'An Airtight Dutch Book'. *Analysis* 59 (1999): 257–265.
Monton, Bradley. 'How to Avoid Maximising Expected Utility'. *Philosophers' Imprint* 19 (2019): 1–25.
Parfit, Derek. *On What Matters, Volume 1* (Oxford, UK: Oxford University Press, 2011).
Podgorski, Abelard. 'Normative Uncertainty and the Dependence Problem'. *Mind* 129 (2020): 43–70.
Ramsey, Frank. *The Foundations of Mathematics and Other Logical Essays* (London: Routledge, 1931).
Regan, Donald. *Utilitarianism and Cooperation* (Oxford, UK: Oxford University Press, 1980).
Robinson, Pamela. 'Is Normative Uncertainty Irrelevant if Your Descriptive Uncertainty Depends on It?' *Pacific Philosophical Quarterly* 103 (2022): 874–899.
Rosenthal, Chelsea. 'What Decision Theory Can't Tell Us About Moral Uncertainty'. *Philosophical Studies* 178 (2021): 3085–3105.
Ross, Jacob. "Acceptance and Practical Reason" (PhD dissertation, Rutgers University, 2006).
Ross, Jacob. 'Rejecting Ethical Deflationism'. *Ethics* 116 (2006): 742–768.
Ross, Jacob. 'Sleeping Beauty, Countable Additivity, and Rational Dilemmas'. *Philosophical Review* 119 (2010): 411–447.
Schroeder, Tim. 'Desire'. In Edward N. Zalta (ed.), *Stanford Encyclopedia of Philosophy* (Summer 2020). Stanford University. https://plato.stanford.edu/archives/sum2020/entries/desire
Sepielli, Andrew. 'What to Do When You Don't Know What to Do'. *Oxford Studies in Metaethics* 4 (2009): 5–28.
Sepielli, Andrew. 'What to Do When You Don't Know What to Do When You Don't Know What to Do . . .'. *Noûs* 48 (2014): 521–544.
Sepielli, Andrew. 'How Moral Uncertaintism Can Be Both True and Interesting'. *Oxford Studies in Normative Ethics 7* (2018): 98–116.
Smith, Nicholas J. J. 'Is Evaluative Compositionality a Requirement of Rationality?' *Mind* 123 (2014): 457–502.
Staffel, Julia. 'Normative Uncertainty and Probabilistic Moral Knowledge'. *Synthese* 198 (2021): 6739–6765.
Sung, Leora. 'Never Just Save the Few'. *Utilitas* 34 (2022): 275–288.
Tarsney, Christian. 'Normative Uncertainty and Social Choice'. *Mind* 128 (2019): 1285–1308,
Tarsney, Christian. 'Vive la Différence? Structural Diversity as a Challenge for Metanormative Theories'. *Ethics* 131 (2021): 151–182.
Weatherson, Brian. 'Review: Moral Uncertainty and Its Consequences'. *Mind* 111 (2002): 693–696.
Weatherson, Brian. 'Running Risks Morally'. *Philosophical Studies* 167 (2014): 141–163.
Wedgwood, Ralph. 'Moral Disagreement and Inexcusable Irrationality'. *American Philosophical Quarterly* 56 (2019): 97–108.
Wilkinson, Hayden. 'In Defence of Fanaticism'. *Ethics* 132 (2022): 445–477.
Williamson, Timothy. *Vagueness* (New York: Routledge, 1994).