Moral uncertainty, expected choiceworthiness, and variance normalization

Harry R. Lloyd

November 12, 2025

Abstract

Many philosophers have recently defended the Maximise Expected Choiceworthiness (MEC) approach to moral uncertainty. Perhaps the most important problem for MEC is the problem of intertheoretic choiceworthiness unit comparisons; and one extant response to this problem is a technique called 'variance normalization.' Although there has been some short, scattered commentary on variance normalization in the philosophical literature, nobody has yet offered a detailed critical discussion of this proposal. This paper fills that lacuna, arguing that there are serious problems which would need to be addressed before variance normalization could be invoked to rescue MEC from scepticism about intertheoretic comparisons.

1 Introduction

Suppose I am uncertain about which moral theory is correct, and that I wish to decide which charity to donate \$100 to. I have 60% credence in a moral theory according to which it is quite important for me to donate all of my money to Oxfam; and I have 40% credence in a moral theory according to which it is morally unimportant which charity I donate to. Under these conditions of moral uncertainty, there is plausibly some sense in which my only *appropriate* option is donating to Oxfam – taking into account the views of both of the moral theories in which I have credence.[1]

[1] 'Appropriate' here is a term of art, that different philosophers explicate in different ways. Perhaps the most popular approach is to understand appropriateness as a kind

In general, which options are appropriate under moral uncertainty? Perhaps the most popular theory of appropriateness is *Maximising Expected Choiceworthiness* (henceforth: MEC). According to MEC:

> (MEC) some option A is appropriate under conditions of moral uncertainty iff choosing A maximises intertheoretic expected choiceworthiness.[2]

The *choiceworthiness* of some option A according to the moral theory T is the strength of the decision maker's moral reasons in favour of choosing A according to T.[3] The *intertheoretic expected choiceworthiness* of some option is a weighted average of its choiceworthinesses according to each of the theories in which the decision maker has credence, where each theory's weight is the decision maker's credence in that theory.

Despite its recent popularity, many philosophers have raised objections to MEC. As I will explain in §2 of this paper, perhaps the most important challenge these critics have posed for advocates of MEC is the problem of *intertheoretic unit comparisons*. And, as I will explain in §3 below, one interesting potential response to this problem is a technique called '*variance normalization*.'

Although there has been some short and scattered commentary on variance normalization in the philosophical literature, no article has yet provided a detailed critical discussion of this proposal. This paper fills that lacuna, and argues that there are serious problems which would need to be addressed before variance normalization could be invoked to rescue MEC from scepticism about intertheoretic comparisons.

Some of the problems that I will discuss here are new to this paper; whereas others have already been hinted at elsewhere, but have not yet received the attention they deserve. Thus far, the problem of 'option individu-

of rationality (e.g. Sepielli 2009; 2014; Bykvist 2014; Geyer 2018; MacAskill, Bykvist and Ord 2020, pp. 20-1). On the other hand, Pittard and Worsnip (2017) defend a metanormative contextualist analysis of appropriateness, according to which the 'ought' of appropriateness is on all fours with the 'ought's of objective and subjective first-order morality. And Olle Risberg (2023) argues that whereas first-order moral theories are truth-apt answers to "the question of what we *ought* to do," theories of appropriateness are conative responses to "the question of what *to* do." I will not assume any particular metaethic of appropriateness in this paper.

[2]Oddie 1994; Lockhart 2000; Sepielli 2009; 2010; Wedgwood 2013; 2017; MacAskill, Bykvist and Ord 2020.

[3]MacAskill, Bykvist and Ord 2020, p. 4.

ation' (see §4 below) has received the greatest amount of philosophical attention. However, I will argue in §§5-7 below that even if the problem of option individuation can be successfully resolved, then variance normalization will still face several even deeper problems. The upshot of this discussion is that it is at present unclear whether some form of variance normalization could ever successfully extricate MEC from the problem of intertheoretic comparisons. This conclusion is bad news for advocates of MEC (who thereby lose this potential response to the problem of intertheoretic comparisons), but good news for advocates of alternative theories of appropriateness such as the bargaining-theoretic approach,[4] or '*My Favourite Theory,*'[5] or '*Maximising Expected Truthlikeness.*'[6]

2 Intertheoretic comparisons

As it is stated in §1, MEC is only applicable in cases where choiceworthiness is *interval-scale measurable* and *intertheoretically unit-comparable* for all of the theories in which one has credence. Choiceworthiness is interval-scale measurable according to some theory T iff for any options A, B, C, and D, none of which are pairwise tied for choiceworthiness according to T, there exists some number k for which it is true and meaningful to say that the difference in choiceworthiness between A and B is k times the size of the difference in choiceworthiness between C and D according to T. Temperature is one example of an interval-scale measurable variable; and °C and °F are two examples of interval scales.

Furthermore, two moral theories T_1 and T_2 are intertheoretically unit-comparable iff for any options A, B, C, and D, none of which are pairwise tied for choiceworthiness according to T_1 or T_2, there exists some number k for which it is true and meaningful to say that the difference in choiceworthiness between A and B according to T_1 is k times the size of the difference in choiceworthiness between C and D according to T_2. The temperature scales °F and °C are unit-comparable with each other: a difference in temperature of 1°C is equivalent to a difference in temperature of 1.8°F.

Trying to calculate intertheoretic expected choiceworthiness is simply incoherent in cases where unit comparisons are impossible. (This would be a

[4] Greaves and Cotton-Barratt 2024; Kaczmarek, Lloyd and Plant 2025; Lloyd 2025b.

[5] Gustafsson and Torpman 2014.

[6] Lloyd 2025a.

bit like trying to calculate the average of 100 acres and 3 bananas.) This is something of a problem for MEC, since many philosophers have been sceptical about the very possibility of intertheoretic unit comparisons.[7] For instance, Edward Gracely suggests that it seems to be "part of the very nature of a moral system that it presents a way of viewing reality, and that the differing visions of different moral systems cannot be directly compared."[8] And Newberry and Ord point out that intertheoretic comparisons seem especially puzzling when the compared moral theories are structurally dissimilar. "For example, it is *prima facie* unclear how to compare how wrong it is to [say] eat pork on a deontological theory, where this is absolutely forbidden, compared to how wrong it is to eat the alternatives under a utilitarian theory, where this is grounded in weighing up the welfare gains and losses."[9]

Furthermore, even in cases where the compared moral theories have significant structural commonalities, and are both interval-scale measurable, intertheoretic unit comparability can still seem implausible. For instance, Brian Hedden considers totalist and averagist utilitarianism.[10] "Suppose we start with a simple proposal – for Totalism we let the value of a world be the total happiness of that world, while for Averagism we let the value of a world be the average level of happiness. Unfortunately," Hedden points out, "this will have the result that for most real-life cases where one can substantially increase population at the cost of decreasing average happiness, [MEC] will recommend doing what Totalism recommends unless the agent is overwhelmingly confident that Averagism is correct."[11]

> We could perhaps improve things by representing Averagism not by the value function that assigns each world its average happiness as its value, but rather by a value function that assigns each world some large multiple of its average happiness as its value. But this proposal is not without its own problems. No matter what value functions we use to represent Averagism and

[7] Hudson 1989; Gracely 1996; Broome 2012, pp. 184-5; Gustafsson and Torpman 2014; Nissan-Rozen 2015; Hedden 2016, §5.2.1; Newberry and Ord 2021; Gustafsson 2022, §5; Kaczmarek, Lloyd and Plant 2025.

[8] Gracely 1996, p. 328

[9] Newberry and Ord 2021, p. 6.

[10] Likewise MacAskill 2014, pp. 93-5. MacAskill also argues that it "seems very difficult to see how there could be any principled way of claiming that there is a unit of value that is shared between utilitarianism and prioritarianism" (MacAskill 2014, pp. 91-3).

[11] Hedden 2016, p. 108.

> Totalism, once we fix on some proposed decrease in average happiness, Averagism will swamp Totalism for smaller population increases while Totalism will swamp Averagism for larger population increases . . . The problem is that representing Totalism and Averagism by particular value functions requires us to choose a point along a continuum of possible population increases where Totalism starts to ouweigh Averagism (for a given reduction in average happiness). And any such choice will seem arbitrary and unmotivated. There is nothing in the moral theories themselves that tells us how to make intertheoretic value comparisons[12]

– and it is difficult to see what other facts could ground or explain these kinds of intertheoretic choiceworthiness comparisons.

3 Variance normalization

As one potential response to these worries about intertheoretic comparisons, several of the most prominent philosophical advocates of MEC have recently suggested *Maximising Expected Variance Normalized Choiceworthiness* (henceforth: MEVNC). MEVNC is an extension of the MEC approach, designed to handle cases in which choiceworthiness *is* interval-scale measurable for all of the theories in which the decision maker has credence, but is *not* intertheoretically unit-comparable across all of these theories.[13] *Variance normalization* is a way of rescaling intertheoretically incomparable choiceworthiness schedules so that the rescaled versions can be compared (for all practical purposes).

Consider some *reference set* of options $\mathcal{O}$, each of which is available in at least one logically possible choice situation. The *variance* of choiceworthiness values over $\mathcal{O}$ according to the moral theory T is a measure of the 'spread' or 'dispersion' of those choiceworthiness values; it is defined as the average squared difference between option choiceworthiness and mean choiceworthiness in $\mathcal{O}$ according to T. Suppose that choiceworthiness variance over $\mathcal{O}$ is

[12] Hedden 2016, pp. 108-9.

[13] MacAskill, Bykvist and Ord 2020; MacAskill, Cotton-Barratt and Ord 2020. In order to extend the MEC approach still further, MacAskill, Bykvist and Ord suggest that we should cardinalise theories that are not interval-scale measurable using the Borda rule voting procedure – see also MacAskill 2016. But for objections to this suggestion, see Tarsney 2019; Carr 2022; Gustafsson 2022; Pivato 2022.

100 according to T_1, but only 20 according to T_2. Then to variance normalize T_1 and T_2's choiceworthiness functions with respect to $\mathcal{O}$ is to multiply each of them by a (different) scaling factor, selected to ensure that the rescaled choiceworthiness functions have the same variances over $\mathcal{O}$. These variance normalized choiceworthiness schedules can then be used to calculate each option's intertheoretic expected variance-normalized choiceworthiness (as we desired).[14]

Variance normalization is one of many possible *statistical normalization* techniques. Each of these possible techniques involves multiplying each choiceworthiness function by a scaling constant selected to ensure that some statistic over $\mathcal{O}$ is equal across choiceworthiness functions. For instance, Ted Lockhart has suggested that choiceworthiness functions should be normalized with respect to their *ranges* over $\mathcal{O}$.[15] (However, this proposal has been convincingly criticised by Andrew Sepielli.[16]) Alternatively, MacAskill, Cotton-Barratt and Ord consider but reject the possibility of normalizing choiceworthiness functions with respect to (i) mean absolute difference (ii) mean absolute deviations around the median in $\mathcal{O}$.[17] Ultimately, MacAskill, Bykvist, Cotton-Barratt and Ord all argue for the superiority of variance normalization over other statistical normalization techniques.[18]

For the sake of simplicity and concreteness, all of my arguments in this

[14]A minor technical problem with variance normalization concerns cases where a decision maker has positive credence in a moral theory T_3 which assigns the same choiceworthiness values to every option in $\mathcal{O}$. Under this assumption, the variance of choiceworthiness values over $\mathcal{O}$ according to T_3 is zero. Hence, it will be impossible to scale down positive-variance theories like T_1 and T_2 to a point where their choiceworthiness variances over $\mathcal{O}$ are equal to T_3's, because any rescaling of T_1 or T_2's choiceworthiness schedules has positive variance over $\mathcal{O}$.

Fortunately, this problem is easy to avoid. To do so, we can amend the variance normalization proposal, stipulating that only theories with positive choiceworthiness variances over $\mathcal{O}$ should have their choiceworthiness schedules variance normalized. Assuming that the relevant reference set $\mathcal{O}$ in any given choice situation contains (at least) all of the options available in this choice situation, it makes no difference what choiceworthinesses a theory like T_3 is normalized to assign to each option in $\mathcal{O}$, since this does not affect which option maximises expected variance normalized choiceworthiness in this situation (MacAskill, Cotton-Barratt and Ord 2020, p. 68).

[15]Lockhart 2000.

[16]Sepielli 2013.

[17]MacAskill, Cotton-Barratt and Ord 2020.

[18]MacAskill, Bykvist and Ord 2020, chapter 4; MacAskill, Cotton-Barratt and Ord 2020; although for some scepticism about the force of these arguments, see Pivato 2022, p. 156.

paper will be framed as objections to variance normalization. However, many of these arguments can quite easily be repurposed into objections against other similar forms of dispersion normalization. If variance normalization faces a host of difficult problems, then there is little reason to think that other forms of dispersion normalization would do much better.

Suppose that in some choice situation, one wants to know the variance normalized choiceworthiness of option A according to the moral theory T. What reference set of options $\mathcal{O}$ should one use to perform the variance normalization? One approach is to use the set of options available in this particular choice situation – call this the *narrow* approach. Another approach is to use something like the set of all options it is possible to confront in any choice situation or other – call this the *broad* approach.[19]

In their 2020 monograph on *Moral Uncertainty*, MacAskill, Bykvist and Ord (henceforth: MB&O) "argue that narrow [variance normalization] is the best approach, though [they] are not confident" in preferring narrowness to broadness.[20] By contrast, in their 2020 paper on 'Statistical normalization methods in interpersonal and intertheoretic comparisons,' the overlapping set of authors MacAskill, Cotton-Barratt and Ord (henceforth: MC&O) argue that "there are two good reasons for preferring the broad formulation."[21] I discuss some objections to narrowness in §§6-7 below; and I discuss an important objection to broadness in §5.[22]

[19]MacAskill, Bykvist and Ord 2020, pp. 101-2; MacAskill, Cotton-Barratt and Ord 2020, p. 68. This characterization of broadness is somewhat ambiguous. One ambiguity concerns the kind of 'possibility' being invoked here: is it logical? Metaphysical? Nomological? Relatedly: are we interested in the set of all possible choice situations for all possible agents; or are we instead interested in the set of choice situations that this particular decision maker could face in the future and (perhaps) could have faced in the past? Different disambiguations in response to these questions will produce different versions of broadness.

[20]MB&O 2020, p. 102.

[21]MC&O 2020, p. 68. More specifically, MC&O (2020, pp. 68-9) argue that (1) "any narrow statistical normalization method generates cyclical recommendations across decision-situations," and, "even more importantly," (2) "it seems clear that some decision-situations are higher stakes for some theories than for others," despite the fact that "narrow accounts do not allow for this."

[22]MB&O (2020, p. 102) object to broadness on the grounds that it would make MEVNC difficult to apply in practice (in reply, cf. MC&O 2020, p. 69). However, I think that this objection understates the case against broadness: I argue in §5 below that there will be many cases in which it will be not just difficult but *impossible* to apply broad MEVNC.

4 Option individuation

One problem that confronts both broad and narrow versions of variance normalization is the problem of *option individuation.*[23] Imagine that according to T_1, option A's choiceworthiness is very close to mean choiceworthiness in the reference set $\mathcal{O}$. By contrast, according to T_2, A's choiceworthiness is far greater than mean choiceworthiness in $\mathcal{O}$. Now suppose that we split option A up into two more fine-grained options. For instance, if option A is pulling the lever in some trolley problem, then we could split A into two options: L, 'pulling the lever with one's left hand;' and R, 'pulling the lever with one's right hand.' Splitting A into L and R will decrease choiceworthiness variance over $\mathcal{O}$ according to T_1, but will increase it according to T_2. Thus, how we individuate our options can change choiceworthiness variance – and hence which options are appropriate according to MEVNC. This is a *prima facie* problem for the variance normalization proposal.

In response to this problem, MB&O suggest that we should redefine the mean and variance of choiceworthiness in $\mathcal{O}$ according to the moral theory T in terms of *weighted* averages, where the weight assigned to any option A is given by some *measure* over the space of possible options. "A measure will define the 'size' of different options, allowing an option to be divided into two smaller options without affecting the variance."[24] If the measure assigns size m to option A, then the sizes of options L and R will together sum to m.

For this response to truly solve the problem of option individuation, MB&O need to say rather more about the nature of their measure. In particular, one is left wondering: what is this measure supposed to supervene on? MB&O introduce a formal solution to their problem without telling us what it corresponds to in the real world. In other words, they posit something whose existence needs to be demonstrated.[25]

In another, earlier paper, MacAskill responds to this challenge by arguing that "it is not important for practical purposes" to have "a theoretical

[23]MB&O 2020, pp. 96-101.

[24]MB&O 2020, p. 99.

[25]As Marcus Pivato puts it (2022, p. 156), MB&O's measure is "a totally arbitrary mathematical gadget, without any motivation or justification; its only role is to make the theory 'work.' [MB&O] do not specify how to choose such a measure, or why one choice is more natural than any other . . . This arbitrary ingredient is the weakest part of their whole framework."

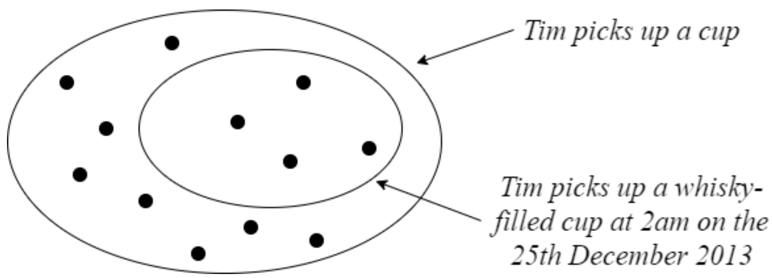

Figure 1: Propositions and possible worlds

grounding for the choice of measure," since "we have an intuitive grasp of the relative 'size' of an option, and the choice of measure is simply a way to make precise that intuitive grasp." According to MacAskill, "it is clear that we have an intuitive notion of the 'size' of possibility space: there is a clear sense in which the proposition [P1] *that Tim picks up a cup* carves out a larger portion of possibility space that the proposition [P2] *that Tim picks up a whisky-filled cup at 2am on the 25th December 2013*." And in order to explain this, "we have to bring in the notion of a measure."[26]

I agree with MacAskill that there is certainly an intuitive sense in which proposition P1 is 'broader' or 'more general' than proposition P2. However, there are alternative ways of explicating this idea which are more ontologically parsimonious than "bring[ing] in the notion of a measure." In particular, we can explicate this idea in terms of the subset relation over propositions construed as sets of possible worlds. Construed in terms of possible worlds, P2 is a subset of P1 (as illustrated in figure 1). This purely ordinal explication does not introduce the cardinal notion of a measure that is required for variance normalization.

Of course, the subset relation is only a partial order over sets of possible worlds. There are many propositions that are neither subsets nor supersets of P1, such as the proposition P3 *that Tim suspends his primary campaign on Fox News on the 12th November 2023*. Thus, my subsethood explication cannot make any sense of the claim that P1 and P3's sizes in possibility space can be compared. Fortunately, this strikes me as entirely appropriate, since it is does not strike me as intuitively plausible to claim that P1 and P3's 'sizes' in possibility space can be meaningfully compared. I cannot make any sense of the question of whether or not P1 is 'larger' than P3. Somewhat

[26]MacAskill 2016, p. 995.

similarly, Christian Tarsney reports (contra MacAskill) that he simply does not find himself with any intuitive grasp of the sizes of options, "or indeed with any clear notion of what the 'size' of an option amounts to."[27]

Marcus Pivato suggests (in passing and without endorsement) that the measure over options could be a Hausdorff measure defined in terms of some similarity metric over the option space.[28] Roughly speaking, the idea here is that P1 is larger than P3 to the extent that the set of ways in which P1 can be made true is more internally heterogeneous than the set of ways in which P3 can be made true. 'Heterogeneity' here implicates some more-primitive notion of similarity over the set of possible worlds.

This proposal would seem to yield strange results in cases where there is a wide range of descriptively different yet normatively equivalent ways of performing some option. For example, imagine a trolley problem in which the lever must be pulled by means of some complicated Rube Goldberg machine. Furthermore, suppose that whereas some Rube Goldberg machines can be activated in a myriad number of dissimilar ways, other Rube Goldberg machines can only be activated in one particular way. Under these circumstances, Pivato's proposal seems to imply that the size of the option to pull the lever should be proportional to the number of different ways of activating the Rube Goldberg machine it is attached to. Intuitively, however, this kind of fact should not affect how much weight this option should be given in determining which actions are appropriate under moral uncertainty,

[27] Tarsney 2019, pp. 1295-6. In an unpublished portion of his dissertation, MacAskill (2014, p. 80) suggests that the measure over options could be defined in terms of (1) "the Lebesgue measure over phase space," (2) "Kolmogrov complexity," (3) "the principle of maximum entropy," or (4) "one's fundamental prior probability distribution over possible worlds." I discuss proposals (3) and (4) later in this section. As for proposals (1) and (2): it is difficult to see what could justify or motivate these proposals, and MacAskill does not offer any suggestions. For a strong objection to (1), see Tarsney 2019, n. 10.

[28] Pivato 2022, p. 156. This proposal is somewhat similar to the measure over options proposed by Jonathan Bennett in his (critical) analysis of the doing-allowing distinction. Bennett (1995, p. 93) suggests that "two propositions about how [an] agent moves at a particular time are to be accorded the same amount of the behaviour space if they are *equally specific*." Although "this will not work with every determinable, e.g. with colours, because for them we have no agreed objective measure of specificity," we do nonetheless "have such measures for space and time, and thus for movement and for specificity of propositions about movement." Bennett goes on to argue that analysing the doing-allowing distinction in these terms suggests that there is no morally salient distinction between doing and allowing. Cf. also Bigelow 1976.

or – hence – in calculating moral theories' variances.[29]

Another response to the problem of specifying the measure over options has been suggested by MC&O. According to MC&O, "an option is a proposition that the decision-maker has the power to make true at a time."[30] Furthermore, every decision maker can be assumed to have some fundamental prior credence distribution over the set of propositions. MC&O suggest that each decision maker should use this credence distribution over propositions as the measure for variance normalization.[31]

For use with broad MEVNC, this choice of measure can be given some *prima facie* plausible motivations (in what follows, I go beyond MC&O's presentation). Suppose, for the sake of illustration, that I am trying to work out which charity it is appropriate for me to donate $100 to. Broad MEVNC implies that first of all, I need to renormalize the choiceworthiness functions of all of the moral theories in which I have positive credence, so that they all share the same variance over a certain broad reference set $\mathcal{O}$.

When I am calculating the variance for each moral theory, how much attention should I pay to each theory's evaluation of the choiceworthiness of my crashing a fighter jet into an alien spaceship on Independence Day? It seems plausible to suppose that I should pay very little attention to each theory's choiceworthiness evaluations of actions like this one that are just 'outlandish.' Surely, MEVNC's appropriateness evaluations should not be strongly affected by theories' choiceworthiness evaluations of actions that a person in my circumstances is extremely unlikely to ever perform in her lifetime?

On the other hand, how much attention should I pay (when I am calculating each theory's choiceworthiness variance) to each theory's evaluation of the choiceworthiness of my ordering a pizza for dinner? It seems plausible to suppose that I should pay quite a bit of attention to each theory's choiceworthiness evaluations of actions like this one that are fairly 'run of the mill.' Surely, MEVNC's appropriateness evaluations should be affected most

[29]One could also argue that Pivato's proposal seems to yield a version of variance normalization that is implausibly sensitive to theories' choiceworthiness evaluations of the option to 'do nothing' in cases where there are many different ways of doing nothing – like smoking a cigarette, wringing one's hands, or chewing some bubblegum (cf. Tarsney 2019, p. 1296).

[30]MC&O 2020, p. 64.

[31]MC&O 2020, pp. 69-70; also MacAskill 2014, p. 80. Similarly, Pivato (2022, p. 156) suggests using the primordial probability from Bolker-Jeffrey decision theory.

strongly by theories' choiceworthiness evaluations of actions that a person in my circumstances is highly likely to perform often in her lifetime?

Another way of expressing this idea might begin with the claim that broad MEVNC's variance calculation procedures should be most sensitive to options that are most salient in the context of my practical decision making. MEVNC is, after all, a proposal for making practical decisions. Furthermore, one might think that salience in this context is well measured by one's credence distribution over options construed as propositions.

MC&O suggest that decision makers should use their ur-prior rather than posterior credence distribution over options as the measure for variance normalization. Using posterior credences "would have the counterintuitive result that the measure over options changes as the decision-maker makes subsequent decision and learns more about the world, thereby changing how different theories are normalized against each other."[32] Using posterior credences also produces a strange result for agents who know that they decide how to act at least partly on the basis of their judgements about appropriateness. For agents like this, facts about which options are appropriate can depend upon the agent's guesses about which options will be appropriate. To the extent that an agent like this has high credence in some option being highly appropriate, she will also have high credence that she will perform this option. In turn, this will influence variance normalization, and hence which options are in fact appropriate. This strikes me as a strange result.

Unfortunately, MC&O's suggestion that decision makers should use their ur-prior credence distribution as the measure also has some strange implications. For instance, if subjective Bayesianism is correct, then two different rational decision makers can have wildly different ur-priors over propositions. Thus, under MC&O's proposed version of MEVNC, which action is most appropriate for a decision maker in some situation can depend upon brute facts about the decision maker's idiosyncratic ur-prior credence distribution. If two moral agents had differing ur-prior distributions, then some option could be appropriate for one but not the other of these agents, even if these two agents have identical posterior credence distributions over moral and empirical theories. One way to avoid this implication would be to adopt a robust form of objective Bayesianism, according to which one particular ur-prior distribution is uniquely rational.[33] However, this is itself a contentious

[32] MC&O 2020, p. 69.

[33] For example, in footnote 27 above, I mentioned MacAskill's suggestion that the mea-

commitment, which many epistemologists will reject.

In this section, I have canvassed several proposals for measures over the space of options. All of these proposals seem to have important drawbacks. Thus, adverting to any of these measures in response to the option individuation problem at best begins to compromise some of the initial appeal of the MEVNC proposal. However, even putting aside the problem of option individuation, there are still several further problems in store for MEVNC.

5 Undefined variances

For the sake of argument, let us now take for granted that choiceworthiness variances are to be calculated in terms of some measure or other over the set of options. Unfortunately, this modification does nothing to guarantee that every moral theory's choiceworthiness variance will be well-defined for broad specifications of the reference set of options $\mathcal{O}$.[34]

For instance, suppose that according to hedonic utilitarianism, the choiceworthiness of any option is equal to the pleasure minus pain that it realizes. Thus, in principle there are no upper or lower limits on choiceworthiness. Any real number n could be the choiceworthiness value of some possible option or other according to utilitarianism.

In itself this fact need not preclude there being a well-defined utilitarian choiceworthiness variance over any given broad reference set. For example, if choiceworthiness values according to utilitarianism were to be distributed over $\mathcal{O}$ in accordance with the standard normal distribution (see figure 2), then choiceworthiness variance over $\mathcal{O}$ according to utilitarianism would just be equal to 1. Under a normal distribution, as we move through choiceworthiness values getting further and further away from zero, the total measure of all possible options having those choiceworthiness values decays asymptotically towards zero. Furthermore, as we move away from zero this decay happens fast enough for variance to be well-defined despite the fact that

sure over options could be defined in terms of the principle of maximum entropy. Some objective Bayesians regard the principle of maximum entropy as a constraint on rational prior credence distributions. However, working out how to apply the principle of maximum entropy in cases where there are infinitely many possible outcomes is itself a difficult problem, that requires further substantive assumptions (see Hájek 2023, §3.1 for an introductory discussion).

[34] Pivato 2022, p. 156; also Gustafsson 2022, p. 462.

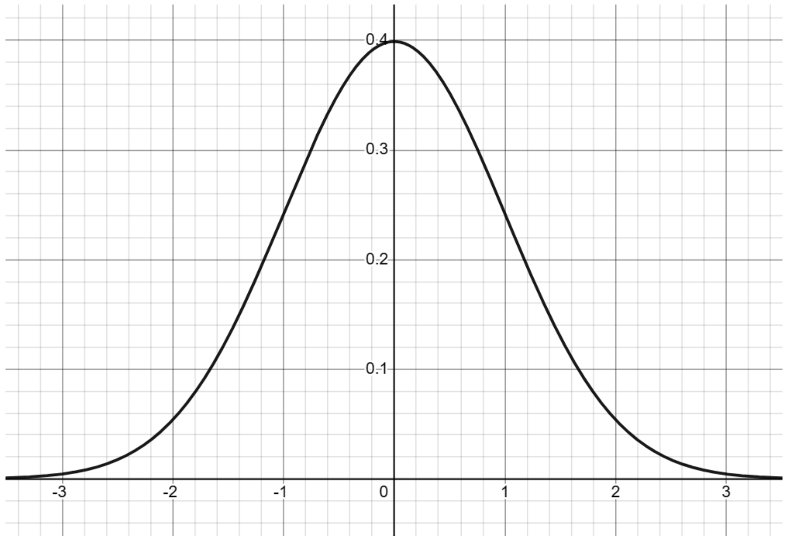

Figure 2: Standard normal distribution

choiceworthiness is unbounded.

Unfortunately, however, there is nothing to guarantee that choiceworthiness values over $\mathcal{O}$ according to utilitarianism will follow a distribution for which the variance is well-defined. For instance, there is nothing to guarantee that these choiceworthiness values will not follow a standard Cauchy distribution – as illustrated in figure 3. Superficially, the standard Cauchy distribution looks quite similar to the standard normal distribution. Just as in the normal distribution, as we move through choiceworthiness values getting further and further away from zero, the total measure of all possible options having those choiceworthiness values decays asymptotically towards zero. Unfortunately, however, as we move away from zero this decay does *not* happen fast enough for variance to be well-defined given that choiceworthiness is unbounded. Visually, the Cauchy distribution is somewhat 'fatter in the tails' than the normal distribution; and this makes all the difference for whether the distribution has a well-defined variance.[35]

[35]Cauchy payoff distributions can also cause problems for standard decision theory, since mean as well as variance is undefined for the Cauchy distribution (as well as for several other similarly pathological payoff distributions) (Alexander 2012). Several philosophers have suggested potential modifications to expected utility maximisation that are designed to sidestep the problem of undefined expectations (for a good overview, see Wilkinson 2025). Unfortunately for advocates of broad MEVNC, none of these modifications have

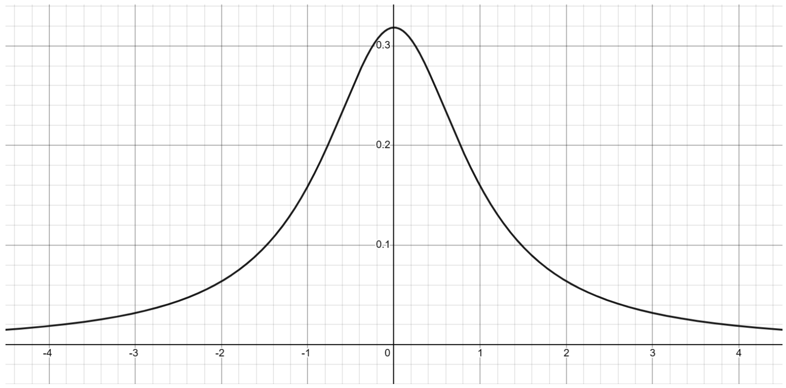

Figure 3: Standard Cauchy distribution

Prima facie, I can see no reason why we should expect choiceworthiness values over any given broad reference set $\mathcal{O}$ to follow a 'well-behaved' distribution according to every moral theory in which one has credence. Indeed, for any antecedently fixed specification of the measure over options which takes nonzero size over an infinite range of options, one can construct a moral theory whose choiceworthiness values over $\mathcal{O}$ will follow an 'ill-behaved' distribution (with an undefined variance) given that measure. (Just make sure that choiceworthiness values according to this moral theory grow in absolute value faster than the measure for those values itself decays.) If some decision maker then has even the slightest credence in this moral theory, she will be unable to use MEVNC since at least one of the moral theories in which she has credence will have an undefined variance.[36]

any obvious analogues that could help us to resolve the undefined variances problem. Partly because of this fact, MEVNC's problem of undefined variances strikes me as likely to be rather less tractable than the problem of undefined expectations that is faced by standard decision theory.

A decision maker who wishes to utilitise standard decision theory will also face the problem of undefined expectations only in a choice situation wherein she has the option to select a lottery whose expected utility is undefined. Whereas, by contrast, a decision maker who wishes to utilise broad MEVNC will face the problem of undefined variances in *every* possible choice situation – provided we grant the modest assumption that this decision maker has some nonzero credence in the claim that choiceworthiness variance is undefined over the reference set $\mathcal{O}$. This is another interesting difference between the problems of undefined expectations and undefined variances.

[36]Actually, *narrow* versions of MEVNC could also suffer from this problem of undefined variances in choice situations for which there is an infinite range of available options (cf.

One potential response to this problem would be to further restrict the scope and ambition of MEVNC. In §1 above, I introduced MEVNC as an extension of the MEC approach designed to handle cases in which choiceworthiness is interval-scale measurable for all of the theories in which one has credence, but is not intertheoretically unit-comparable across all of these theories. In response to the problem introduced in this section, one could restrict the scope of broad MEVNC still further, framing it as only designed to handle cases in which choiceworthiness variance is well-defined for all of the theories in which one has positive credence.

This would be a disappointing concession. Firstly, the problem of undefined variances has the potential to arise even for familiar moral theories like utilitarianism. The appeal of broad MEVNC would be significantly diminished if it could not even be relied upon to handle familiar moral theories like this. Secondly, restricting the scope of broad MEVNC in this way simply raises the question of how advocates of MEC *should* handle theories whose choiceworthiness variances over $\mathcal{O}$ are not well-defined. One approach here would be to ignore these theories' interval-scale measurements, treating them on a par with merely ordinal theories.[37] However, it would be undesirable to ignore interval-scale information in this manner.[38] It would be more desirable to try to develop a decision procedure that is sensitive to the interval-scale information involved in these theories without requiring variance normalization.[39] However, if such a procedure could be developed, then it is not clear why one would not just wholesale replace broad MEVNC with this new decision procedure.

6 Option addition

In the remainder of this paper, I will turn my attentions specifically to the prospects for *narrow* versions of MEVNC. I will argue that narrow MEVNC faces several important problems.

The first of these is the problem of *option addition*. To illustrate this

§7 below). However, to press this objection against narrow MEVNC would require me to introduce several further complications which I wish to prescind from discussing here.

[37] See footnote 13 above.

[38] Tarsney 2021, §VI.

[39] Cf. Greaves and Cotton-Barratt 2024; Kaczmarek, Lloyd and Plant 2025; Lloyd 2025a; 2025b.

Choiceworthiness (to two decimal places)	A	B	C
T_1	0	-1	-2
T_2	-2	0	-1
T_3	-0.8	-1.99	0
Expectation	-0.93	-0.99	-1

Table 1: Initial trolley problem

problem, consider a scenario in which a runaway trolley is headed toward five people who are tied to the tracks. The moral agent in this scenario has three options:

> (A) Pull a lever to redirect the trolley onto a side track where only two people are trapped, killing them but saving the five.
>
> (B) Do nothing, allowing the five to die.
>
> (C) Quickly push a heavyset man wearing earmuffs into the path of the trolley, killing him but saving the five.

Furthermore, suppose that the moral agent has $33.\dot{3}\%$ credence each in three moral theories, T_1, T_2 and T_3. Table 1 represents T_1 through T_3's choiceworthiness schedules for this trolley problem.

Finally, suppose for the sake of simplicity that options A, B, and C all have equal measure. Under this assumption, T_1 through T_3's choiceworthiness schedules all have the same variance over $\{A, B, C\}$, and so narrow variance normalization would not alter these choiceworthiness schedules. Hence, as table 1 shows, option A uniquely maximises expected narrow variance normalized choiceworthiness.

However, now suppose that before our moral agent can make her choice, I remove the heavyset man's earmuffs. This creates one additional option:

> (D) Make a loud noise that the heavyset man will mistake for something dangerous, causing him to jump into the path of the trolley. (Suppose that there is no time to try to persuade him to jump in front of the trolley of his own volition.)

Suppose, furthermore, that the moral theories T_1, T_2 and T_3 all regard option D as a mere 'duplicate' of option C – having the choiceworthiness schedules specified in table 2.

Choiceworthiness (to two decimal places)	A	B	C	D
T_1	0	-1	-2	-2
T_2	-2	0	-1	-1
T_3	-0.8	-1.99	0	0

Table 2: Modified trolley problem, unnormalized

Normalized choiceworthiness (to 2 d.p.)	A	B	C	D
T_1	0	-1	-2	-2
T_2	-2.34	0	-1.17	-1.17
T_3	-0.81	-2.03	0	0
Expectation	-1.05	-1.01	-1.06	-1.06

Table 3: Modified trolley problem, variance normalized

On any plausible measure over options, D will have nonzero measure. Hence, adding D into the reference set has the potential to significantly change each moral theory's choiceworthiness variance over that set. This in turn could quite easily change whether it is appropriate according to narrow MEVNC for the agent to kill the heavyset man, pull the lever, or do nothing in this choice situation. Yet this implication strikes me as intuitively implausible, given that the new option D is merely a 'duplicate' of the previously available option C.[40] This is the problem of option addition.

Suppose for the sake of illustration that options A, B, C and D all have equal measure. Under that assumption, table 3 represents one way of variance normalizing T_1 through T_3's choiceworthiness schedules with respect to $\{A, B, C, D\}$. Hence, as table 3 shows, options B now uniquely maximises expected narrow variance normalized choiceworthiness.

Therefore, according to narrow MEVNC, adding option D into the mix changes which option is uniquely appropriate from A to B, despite the fact that D is merely a duplicate of the already available option C. This result strikes me as intuitively implausible. Why should adding option D into the mix improve B's standing, relative to A?

Note that this implausible feature of narrow MEVNC does not depend on our simplifying assumption that options A, B, C and D all have equal measure. So long as option D has nontrivial measure, adding D into the reference set $\mathcal{O}$ has the potentially to nontrivially change each moral theory's

[40]Likewise Tarsney 2019, pp. 1293-5; Gustafsson 2022, n. 22.

choiceworthiness variance over $\mathcal{O}$.

Clearly, it will not do to suggest that we could solve the problem of option addition by assigning option D zero measure. For consider a scenario in which only options A, B, and D are available (suppose that the heavyset man is too heavy for the agent to push). In this version of the scenario, it is intuitively plausible that when the moral agent is calculating her moral theories' choiceworthiness variances, she should pay attention to option D as well as to options A and B. It would be both *ad hoc* and intuitively implausible to suggest otherwise that the measure should be zero for option D.

Alternatively, one might think that we could solve this problem by shifting from a measure over options to a measure over *normatively-equivalent classes* of options. If our moral agent is certain that options C and D differ only in morally irrelevant features, then she could claim that C and D's total measure relative to the option set $\{A, B, C, D\}$ should be equal to C's measure relative to the option set $\{A, B, C\}$.

However, one important problem with this proposal is that narrow MEVNC can have implausible implications even in option addition cases where the decision maker is certain that the additional option is normatively inequivalent to all of the original options. For instance, suppose that according to some of the moral theories in which the agent has credence, (D) deceiving the man into thinking that there is something dangerous nearby is very marginally worse than (C) physically battering the heavyset man. On the other hand, according to the other moral theories in which the agent has credence, (C) physically battering the heavyset man is very marginally worse than (D) deceiving him. Thus, although our agent is certain that C and D differ only very marginally in normatively relevant features, she is also certain that C and D are normatively inequivalent. We might say that in normative terms, C and D are 'near duplicates.'

Under these assumptions, it strikes me as plausible to suppose that adding D to the original three-option trolley scenario should make very little difference to the variances of the moral theories in which the agent has credence, and hence should be very unlikely to change whether it is appropriate according to narrow MEVNC for the agent to kill the heavyset man, pull the lever, or do nothing. Even assuming that our measure is defined over normatively-equivalent classes of options rather than over options themselves, the only way to achieve these desired results is for D's measure to be drastically lower than the measures of A, B, and C. However, this is no more plausible than the suggestion that D's measure should be zero. (Once again, consider a scenario in which only options A, B, and D are available. In this version of the scenario, it is intuitively plausible that when the agent is calculating the moral theories' choiceworthiness variances, she should not pay drastically more attention to options A and B than she pays to option D. It would be both *ad hoc* and intuitively implausible to suggest otherwise that the measure should be near zero.)

All things considered, then, I struggle to see how narrow MEVNC could avoid the problem of option addition. In formal terms, avoiding this problem would at the very least require us to relativize the measure over options with respect to option sets. Supposing our moral agent is certain that C and D are duplicates in every normatively relevant respect, then what we *want* is for the sum of C and D's measures relative to the option set $\{A, B, C, D\}$ to be equal to C's measure alone relative to the option set $\{A, B, C\}$. And if C and D are near-duplicates for our moral agent, then we probably want the sum of C and D's measures relative to the option set $\{A, B, C, D\}$ to be only slightly larger than C's measure alone relative to the option set $\{A, B, C\}$.

However, although it's easy to say that we want those results in this toy example, it is much harder to say what the rules should be in general. So the ball is now back in the MEVNC advocates' court. The burden of proof is on them to show how the option addition problem can be satisfactorily avoided.

7 Scenario individuation

Another problem with narrow MEVNC is the problem of *scenario individuation*. Once again, imagine a trolley problem in which the moral agent has three options. A runaway trolley is headed toward five people who are tied to the tracks. The agent could push one heavyset man off a bridge onto the tracks, killing him but saving the five. Alternatively, the moral agent could instead run over to a lever. Pushing this lever will redirect the trolley onto a side track where only two people are trapped, killing them, but saving the five. Her final option is to do nothing. Suppose for the sake of simplicity that all three of these options have equal measure.

Furthermore, imagine that our moral agent has 40% credence in some moral theory T_4, and 60% credence in the moral theory T_5. Choiceworthiness is interval-scale measurable according to both of these theories, but is unit-incomparable between them. Table 4 represents T_4 and T_5's choiceworthiness

Choiceworthiness	Killing the heavyset man	Switching the tracks	Doing nothing
T_4	0	-1	-5
T_5	-5	-1	0
Expectation	-3	-1	-2

Table 4: Three-option trolley problem

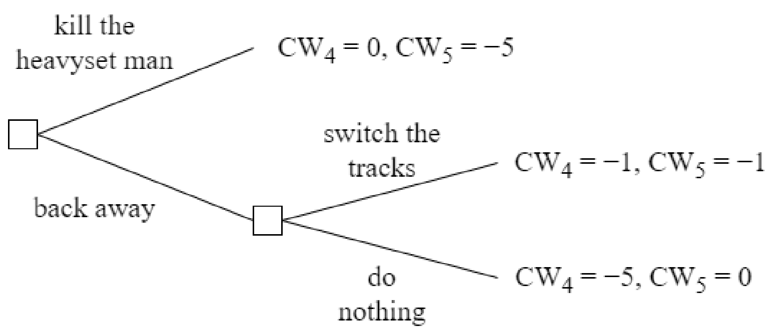

Figure 4: Two-stage construal

schedules for this trolley problem.

Suppose we say that this trolley problem should be understood as a single choice situation. It that case, T_4 and T_5's narrow variance normalized choiceworthiness schedules will be identical to the unnormalized schedules given in table 4. Thus, switching the tracks uniquely maximises expected narrow variance normalized choiceworthiness.

On the other hand, now suppose we say that this trolley problem should be understood as a sequence of two choice situations. If the moral agent chooses to back away from the heavyset man, then after that she will face a second choice, between switching the tracks and doing nothing (as illustrated in figure 4). T_4 and T_5's narrow variance normalized choiceworthiness schedules for the second of these two choice situations are given in table 5. As this table also indicates, doing nothing uniquely maximises expected narrow variance normalized choiceworthiness in this second choice situation.

What about the choice between killing the heavyset man and backing away from him? Each moral theory's choiceworthiness evaluation of the action of 'backing away' might well depend upon what the moral agent would subsequently do immediately after backing away.

Normalized choiceworthiness	Switching the tracks	Doing nothing
T_4	-1	-5
T_5	-4	0
Expectation	-2.8	-2

Table 5: Second trolley sub-problem

Normalized choiceworthiness	Killing the heavyset man	Backing away
T_4	0	-5
T_5	-5	0
Expectation	-3	-2

Table 6: First trolley sub-problem

Let us suppose that our moral agent knows she will always maximise expected variance normalized choiceworthiness in any choice situation she encounters. Furthermore, suppose that according to both T_4 and T_5, the choiceworthiness of 'backing away' in the first choice situation is equal to the choiceworthiness of whatever the agent will in fact do next. Hence, according to both T_4 and T_5, backing away has the same choiceworthiness as 'doing nothing' in the second choice situation. So in the first choice situation, T_4 and T_5's choiceworthiness schedules are given in table 6. Moreover, narrow variance normalization will not alter these choiceworthiness schedules. Backing away from the heavyset man uniquely maximises expected narrow variance normalized choiceworthiness in the first choice situation.

Thus, if we say that this trolley problem is a single choice situation, then narrow MEVNC implies that switching the tracks is the only appropriate option. On the other hand, if we say that this trolley problem is a sequence of two choice situations, then narrow MEVNC implies that doing nothing is the only appropriate course of action. Hence, narrow MEVNC's verdicts are sensitive to how we individuate choice situations.

This result has particularly important implications for the debate between advocates of MEC and philosophers who have defended *bargaining-theoretic* responses to moral uncertainty.[41] In their recent paper on bargaining-theoretic approaches, Greaves and Cotton-Barratt point out that these approaches face problems of scenario individuation.[42] However, Greaves and Cotton-Barratt

[41] Greaves and Cotton-Barratt 2024; Kaczmarek, Lloyd and Plant 2025; Lloyd 2025b.

[42] Greaves and Cotton-Barratt 2024, §10.

do not mention that narrow MEVNC also suffers from this problem. To my knowledge, nobody has yet discussed scenario individuation as a problem for narrow variance normalization.

Greaves and Cotton-Barratt conclude that "whilst the bargaining-theoretic approach is not obviously superior to an MEC approach ... neither is it clearly inferior."[43] One of their key reasons for reaching this ambivalent conclusion is their observation that bargaining-theoretic approaches face problems of scenario individuation.[44] However, recognising that narrow MEVNC also suffers from the scenario individuation problem arguably shifts the balance of advantages somewhat in favour of bargaining approaches. In response to objections concerning scenario individuation, bargaining theorists may respond *tu quoque* that the narrow MEVNC development of the MEC approach is vulnerable to exactly the same objections.

It is most natural *prima facie* to construe my three-option trolley problem as one single choice situation, thus individuating this trolley scenario coarse-grainedly. Under that construal MEVNC implies that switching the tracks is the uniquely appropriate course of action in this trolley problem. Intuitively, this strikes me as the right result.

However, the real trouble for MEVNC lies in finding and defending a universally-applicable rule for scenario individuation. The literature on moral uncertainty has typically treated *individuation problems* like this as serious disadvantages for theories of appropriateness that are subject to these problems. For instance, MB&O themselves regard the problem of *theory individuation* as a decisive reason to reject the '*My Favourite Theory*' approach to moral uncertainty.[45] Somewhat more tentatively, I suggest that finding a plausible way to avoid the problem of scenario individuation is at least an urgent challenge for advocates of narrow MEVNC.

More precisely, the challenge is to find a principled alternative to the maximally 'grand-world' approach to scenario individuation. According to this *maximalist* approach, one's choice situation at any given point in time when one has not yet decided what to do is a choice about how to live the entirety of the rest of one's life. Each option is a (contingent) plan for the rest of one's life, covering every decision node that one might arrive at whilst implementing this complete plan.[46] In other words, *each* maximalist

---

[43] Greaves and Cotton-Barratt 2024, §11

[44] Greaves and Cotton-Barratt 2024, §11.

[45] MB&O 2020, pp. 41-4.

[46] Cf. Lockhart 2000, chapter 8.

- First, back away from the heavyset man.
- Then, redirect the trolley by switching the tracks.
- After that, notify emergency services.
- After that,
  - *if* the police arrive at the scene first, then describe to them what happened. Then after that,
    - *if* ... then ...
    - and *if* ... then ...
  - and *if* an ambulance arrives at the scene first, then ... Then after that, ...
  - and *if* ..., then ...
  - ... and so on and so forth.

Figure 5: A maximal option

option will have something like the structure illustrated in figure 5. As figure 5 should make clear, the maximalist conception of options and of choice situations is very far removed from our intuitive understanding of these concepts.

Maximalism also shares several problems with broadness. Firstly, maximalism makes MEVNC's theory of appropriateness difficult to apply in practice.[47] In the words of Savage, something like maximalism

> demands that one envisage every conceivable policy for the government of his whole life (at least from now on) in its most minute details, in the light of the vast number of unknown states of the world, and decide here and now on one policy. This is utterly ridiculous, ... because the task implied in making such a decision is not even remotely resembled by human possibility.[48]

Secondly, and I think more importantly, maximalism leaves open the

[47] Cf. MB&O's (2020, p. 102) objections to broadness.

[48] Savage 1972, p. 16.

possibility that some moral theories' variances will be undefined over the reference set (recall §4 above). The range of possible maximalist life plans is almost limitless; I can see no reason why there should be bounds, for instance, upon utilitarian choiceworthiness over the maximalist set of life plans. Thus, there is nothing to guarantee that utilitarianism's choiceworthiness variance will be well-defined under maximalism.

These problems tell against a maximalist theory of scenario individuation. The challenge for defenders of narrow MEVNC is to find an alternative theory that is coarse-grained enough to have plausible implications in cases like my trolley problem, but that does not go as far as maximalism. In other words: the challenge is to find and defend some 'stopping point' along the path from fine-grainedness to maximalism. Greaves and Cotton-Barratt worry that "any such choice (short of the impractical maximally grand-world model) seems arbitrary."[49]

Of course, nothing I have written in this section rules out there being some principled, non-"arbitrary" theory of scenario individuation. However, unless and until such an approach can be found, scenario individuation will be a significant unresolved problem for narrow MEVNC.

8 Conclusion

In this paper, I have discussed several problems for both the broad and narrow versions of MEVNC, and I have argued that both versions of MEVNC face significant unresolved problems. Of course, none of my arguments in this paper conclusively prove that the variance normalization idea is totally unsalvageable – there might be ways to modify MEVNC so as to avoid all of the problems that I have discussed here. Nonetheless, it is at best unclear whether some form of variance normalization could ever successfully extricate MEC from the problem of intertheoretic unit comparisons. Perhaps this conclusion should motivate us to more seriously consider some of the budding alternatives to MEC, like bargaining-theoretic approaches[50] or 'Maximising Expected Truhthlikeness.'[51] Given the difficult problems with variance normalization, advocates of MEC will likely have to rely upon other

[49]Greaves and Cotton-Barratt 2024, §11.
[50]Greaves and Cotton-Barratt 2024; Kaczmarek, Lloyd and Plant 2025; Lloyd 2025b.
[51]Lloyd 2025a.

ideas if they wish to avoid the problem of intertheoretic comparisons.[52]

Acknowledgements

For helpful comments and conversations, I wish to thank Paul Forrester, Dan Greco, Shelly Kagan, Andreas Mogensen, Marcus Pivato, and Christian Tarsney.

References

Alexander, J. McKenzie. 2012. Decision theory meets the witch of Agnesi. *Journal of Philosophy*, 109.12, 712-27.

Bennett, Jonathan. 1995. *The Act Itself* (Oxford: Clarendon Press).

Bigelow, John C. 1976. Possible worlds foundations for probability. *Journal of Philosophical Logic*, 5.3, 299-320.

Broome, John. 2012. *Climate Matters: Ethics in a Warming World* (New York, NY: W. W. Norton).

Bykvist, Krister. 2014. Evaluative uncertainty, environmental ethics, and consequentialism. Pp. 122-35 in Avram Hiller, Ramona Ilea and Leonard Kahn (eds), *Consequentialism and Environmental Ethics* (New York, NY: Routledge).

Carr, Jennifer Rose. 2022. The hard problem of intertheoretic utility comparisons. *Philosophical Studies*, 179.4, 1401-27.

Ecoffet, Adrien and Lehman, Joel. 2021. Reinforcement learning under moral uncertainty. *Proceedings of the 38th International Conference on Machine Learning*, Proceedings of Machine Learning Research 139, 2926-36.

Geyer, Jay. 2018. Moral uncertainty and moral culpability. *Utilitas*, 30.4, 399-416.

Gracely, Edward J. 1996. On the comparability of judgments made by different ethical theories. *Metaphilosophy*, 27.3, 327-32.

Greaves, Hilary and Cotton-Barratt, Owen. 2024. A bargaining-theoretic approach to moral uncertainty. *Journal of Moral Philosophy*, 21.1-2, 127-69.

Gustafsson, Johan E. 2022. Second thoughts about my favourite theory. *Pacific Philosophical Quarterly*, 103.3, 448-70.

[52]Cf., for example, Riedener 2021 and Makins 2025.

Gustafsson, Johan E. and Torpman, Olle. 2014. In defence of My Favourite Theory. *Pacific Philosophical Quarterly*, 95.2, 159-74.

Hájek, Alan. 2023. Interpretations of probability. In E. N. Zalta (ed), *The Stanford Encyclopedia of Philosophy*, Winter 2023.

Hedden, Brian. 2016. Does MITE make right? On decision-making under normative uncertainty. Pp. 120-28 in Russ Shafer-Landau (ed), *Oxford Studies in Metaethics*, volume 11 (Oxford: Oxford University Press).

Hudson, James L. 1989. Subjectivization in ethics. *American Philosophical Quarterly*, 26.3, 221-9.

Kaczmarek, Patrick, Lloyd, Harry R. and Plant, Michael. 2025. Moral uncertainty, proportionality, and bargaining. *Ergo*, 12.44, 1142-71.

Lloyd, Harry R. 2025a. Moral uncertainty and expected truthlikeness. *Synthese*, 206.265, 1-32.

Lloyd, Harry R. 2025b. *Moral Uncertainty and Bargaining.* PhD dissertation. New Haven, CT: Department of Philosophy, Yale University.

Lockhart, Ted. 2000. *Moral Uncertainty and Its Consequences* (Oxford: Oxford University Press).

MacAskill, William. 2014. *Normative Uncertainty.* DPhil dissertation. Oxford: Department of Philosophy, University of Oxford.

MacAskill, William. 2016. Normative uncertainty as a voting problem. *Mind*, 125.500, 967-1004.

MacAskill, William, Bykvist, Krister and Ord, Toby. 2020. *Moral Uncertainty* (Oxford: Oxford University Press).

MacAskill, William, Cotton-Barratt, Owen and Ord, Toby. 2020. Statistical normalization methods in interpersonal and intertheoretic comparisons. *Journal of Philosophy*, 117.2, 61-95.

Makins, Nicholas. 2025. Suppositional desires and rational choice under moral uncertainty. *Ergo*, 12.10, 247-66.

Newberry, Toby and Ord, Toby. 2021. The parliamentary approach to moral uncertainty. Future of Humanity Institute, technical report 2021-2.

Nissan-Rozen, Ittay. 2015. Against moral hedging. *Economics and Philosophy*, 31.3, 349-69.

Oddie, Graham. 1994. Moral uncertainty and human embryo experimentation. Pp. 144-61 in K. W. M. Fulford, Grant Gillett, and Janet Martin Soskice (eds), *Medicine and Moral Reasoning* (Cambridge: Cambridge University Press).

Pittard, John and Worsnip, Alex. 2017. Metanormative contextualism and normative uncertainty. *Mind*, 126.501, 155-93.

Pivato, Marcus. 2022. Review of *Moral Uncertainty*. *Economics and Philosophy*, 38.1, 152-8.

Riedener, Stefan. 2021. *Uncertain Values: An Axiomatic Approach to Moral Uncertainty* (Berlin: de Gruyter).

Risberg, Olle. 2023. Ethics and the question of what to do. *Journal of Ethics and Social Philosophy*, 25.2, 376-412.

Savage, Leonard J. 1972. *The Foundations of Statistics*, 2nd edition (New York: Dover).

Sepielli, Andrew. 2009. What to do when you don't know what to do. Pp. 5-28 in Russ Shafer-Landau (ed), *Oxford Studies in Metaethics*, volume 4 (Oxford: Oxford University Press).

Sepielli, Andrew. 2010. *'Along an Imperfectly Lighted Path': Practical Rationality and Normative Uncertainty.* PhD dissertation. New Brunswick, NJ: Department of Philosophy, Rutgers University.

Sepielli, Andrew. 2013. Moral uncertainty and the principle of equity among moral theories. *Philosophy and Phenomenological Research*, 86.3, 580-9.

Sepielli, Andrew. 2014. What to do when you don't know what to do when you don't know what to do . . . . *Noûs*, 48.3, 521-44.

Tarsney, Christian J. 2019. Normative uncertainty and social choice. *Mind*, 128.512, 1285-308.

Tarsney, Christian J. 2021. Vive la différence? Structural diversity as a challenge for metanormative theories. *Ethics*, 131.2, 151-82.

Wedgwood, Ralph. 2013. *Akrasia* and uncertainty. *Organon F*, 20.4, 484-506.

Wedgwood, Ralph. 2017. Must rational intentions maximise utility? *Philosophical Explorations*, 20.S2, 73-92.

Wilkinson, Hayden. 2025. Flummoxing expectations. *Noûs*, 59.3, 700-28.