

# 5

## Intertheoretic Comparisons of Choice-Worthiness

### Introduction

So far, we have given an account of how to make decisions in the face of moral uncertainty that can be applied even when some of the theories in which one has credence are not comparable with each other. This raises the question: how often are differences of choice-worthiness comparable across theories? (In this book, we only consider the issue of intertheoretic comparisons of choice-worthiness differences. There is a separate question of whether *levels* of choice-worthiness are comparable across theories. However, as noted in the introduction, we do not discuss level-comparability of choice-worthiness in this book. When we use the term *comparable* in every instance we're referring to comparability of differences of choice-worthiness, not level-comparability.)

Three distinct lines of argument suggest that intertheoretic comparisons of choice-worthiness differences are impossible or, if possible, are bound to lead to implausible normative results.

First, the *appeal to cases* argument. In many cases there seems to be no intuitive way in which to compare two moral theories. As noted in Chapter 4, even for theories as similar as utilitarianism and prioritarianism, there appears to be no principled way of determining whether prioritarianism is agreeing with utilitarianism about the value of wellbeing-increases for the well-off but claiming that those for the badly-off get extra weight, or whether it is agreeing for the badly-off and claiming that wellbeing-increases for the well-off matter less (or some third option).

Second, the *swamping* argument. Even in some cases where there does seem to be a 'natural' way to compare the two theories, this natural comparison quickly leads to implausible results, causing one theory to 'swamp' the other in the expected choice-worthiness calculation. Brian Hedden makes this argument with respect to Average and Total Utilitarianism

(where he considers equating a unit of total wellbeing with a unit of average wellbeing):<sup>1</sup>

Suppose that the agent has the choice of increasing the world's population from 6 billion to 24 billion people at the cost of halving the average happiness level... maximizing intertheoretic expectation will recommend that the agent implement the population-increasing policy (i.e. doing what Totalism recommends) unless she is over 99.9999999916% confident that Averagism is right. But this seems crazy.

After considering different ways of giving the two theories a common unit (or 'normalizing' those theories), he concludes that the problem is unresolvable: 'No matter what value functions we use to represent Averagism and Totalism, once we fix on proposed decrease in average happiness, Averagism will swamp Totalism for smaller population increases while Totalism will swamp Averagism for larger population increases.'

Again, however, the fact that these are both such similar theories should make us worried. If we can't make plausible choice-worthiness comparisons across two very similar versions of utilitarianism, what hope do we have to make comparisons across very different sorts of theory, such as utilitarianism and virtue ethics?<sup>2</sup>

The third and most general argument is the *arbitrary unit* argument.<sup>3</sup> The natural way of understanding *intratheoretic* comparisons of choice-worthiness differences, so the argument goes, is in terms of how a theory orders prospects under empirical uncertainty: that what it means for the difference in choiceworthiness between *A* and *B* (where *A* is more choiceworthy

<sup>1</sup> Brian Hedden, 'Does MITE Make Right? On Decision-Making under Normative Uncertainty', *Oxford Studies in Metaethics*, vol. 11 (2016), p. 108. As he notes, this argument can also be found in William MacAskill, 'Normative Uncertainty', DPhil Thesis, University of Oxford, 2014; the progenitor of the case is Toby Ord. The example of average and total utilitarianism is also given by John Broome, *Climate Matters: Ethics in a Warming World*, New York: W. W. Norton, 2012, p. 185 as part of an assertion that intertheoretic comparisons are almost always impossible, though he doesn't make the *swamping* argument. These cases are considered in depth in Hilary Greaves and Toby Ord, 'Moral Uncertainty about Population Axiology', *Journal of Ethics and Social Philosophy*, vol. 12, no. 2 (November 2017), pp. 135–67. The authors are inclined to accept the *swamping* conclusion as a *modus ponens*.

<sup>2</sup> A different response would be to say that the problem is not with the intertheoretic comparison, but with maximizing expected choice-worthiness (which allows some theories to swamp others). We discuss a related issue in our section on 'fanaticism' in Chapter 6.

<sup>3</sup> This argument is made in Ittay Nissan-Rozen, 'Against Moral Hedging', *Economics and Philosophy*, vol. 31, no. 3 (November 2015), pp. 349–69.

than  $B$ ) to be equally as large as the difference in choice-worthiness between  $B$  and  $C$  (where  $B$  is more choiceworthy than  $C$ ) is that a guarantee of  $B$  is equally as good as the prospect of a 50/50 chance of either  $A$  or  $B$ . More precisely, if a theory orders all possible prospects in terms of their choice-worthiness and satisfies the von Neumann–Morgenstern axioms, then it can be represented as maximizing expected choice-worthiness, where choice-worthiness can be represented on an interval scale.<sup>4</sup> However, for each theory, the numerical representation of its choice-worthiness ordering is unique only up to a positive affine transformation: the unit is arbitrary. So if there is nothing more to choice-worthiness than an individual theory's choice-worthiness ordering over prospects (and this is an assumption we will return to later in the chapter), then the choice of unit is arbitrary for the representation of the choice-worthiness ordering of each moral theory, and it's meaningless to say that one unit of choice-worthiness, on one theory, is greater than, smaller than, or equal to one unit of choice-worthiness on another theory.

These worries have given rise to three classes of response. The *sceptics* argue that choice-worthiness differences are either always incomparable across theories<sup>5</sup> or are almost always incomparable across theories.<sup>6</sup>

The *structuralists* claim that intertheoretic comparisons are often possible, and that intertheoretic comparisons should be made only with reference to structural features of the theories' qualitative choice-worthiness relation (such as the choice-worthiness of the best option and worst option) or mathematical features of its numerical representation (such as the mean, sum, or spread of choice-worthiness). One might believe, for example, that variance voting is not merely the best way to act if theories are incomparable, but that it is the correct way to actually make intertheoretic comparisons; this would be a structural account.<sup>7</sup> The non-structuralists deny

<sup>4</sup> Von Neumann and Morgenstern, *Theory of Games and Economic Behavior*. Note, though, that their concern is preference-orderings rather than with choice-worthiness orderings. For discussion of von Neumann and Morgenstern's result to moral theories see Broome, *Weighing Goods*.

<sup>5</sup> Hudson, 'Subjectivization in Ethics'; Gracely, 'On the Noncomparability of Judgments Made by Different Ethical Theories'; Gustafsson and Torpman, 'In Defence of My Favourite Theory' and perhaps Hedden, 'Does MITE Make Right?'

<sup>6</sup> John Broome, 'The Most Important Thing about Climate Change', in Jonathan Boston, Andrew Bradstock, and David L. Eng (eds), *Public Policy: Why Ethics Matters*, Acton, ACT: Australia National University E Press, 2010, pp. 101–16; Broome, *Climate Matters*, p. 122.

<sup>7</sup> One interpretation of Lockhart, *Moral Uncertainty and Its Consequences* is that he's a structuralist. Other structural accounts are suggested by Sepielli, 'Normative Uncertainty for Non-Cognitivists'.

structuralism, and take into account something more than just positional features of theories' choice-worthiness functions.<sup>8</sup>

In this chapter we will argue against both scepticism and structuralism. We believe that understanding why both scepticism and structuralism are false helps us to understand the scope of possibilities within non-structural accounts; our arguments will therefore also provide some positive reasons in favour of non-structuralism, and some suggestions of possible non-structural accounts.

## I. Against Scepticism

The first two arguments in favour of scepticism appealed to specific cases, where either there were no intuitive comparisons to be made, or where the natural comparison would lead to swamping, which was taken to be an implausible result.

However, one can only draw a limited conclusion by appealing to specific cases. At most, one can show that *sometimes* intertheoretic comparisons do not hold between two theories. One cannot, thereby, show that they (almost) *never* hold between two theories, or that they are impossible. Usually, the *appeal to cases* and *swamping* arguments have been made in the context of arguing against MEC. A presupposition has been that if intertheoretic comparisons of choice-worthiness differences are *sometimes* impossible, then MEC cannot be a perfectly general account of what to do under normative uncertainty. But that presupposition is false; as we have seen in the previous chapters, by using Borda and variance voting we can apply a modified form of MEC even in conditions of intertheoretic comparability and even in conditions of merely ordinal theories.

Moreover, there are also many cases where two different moral views intuitively *do* seem comparable. We describe three classes of cases.

The first class of cases is the most compelling cases of MEC-style reasoning, where the two moral viewpoints differ with respect to only one moral issue. Consider, for example, the following statements.<sup>9</sup>

<sup>8</sup> Non-structural accounts are suggested by Ross, 'Rejecting Ethical Deflationism', pp. 763–4 and Sepielli, "Along an Imperfectly Lighted Path".

<sup>9</sup> In each example statement we give in this section, we will use the natural English locution to make the intertheoretic comparison. However, strictly speaking we should consider the natural English as shorthand. So, when we imagine someone saying, 'If animals have rights in the way that humans do, then killing animals is a much more severe wrongdoing than if they don't

If animals have rights in the way that humans do, then killing animals is a much more severe wrongdoing than if they don't.

If Singer is right about our duties to the poor, then our obligation to give to development charities is much stronger than if he's wrong.

These are cases where we're not really comparing two different complete theories, considered in the abstract. We're comparing two different moral views that differ with respect to just one moral issue. In these cases, the intertheoretic comparison seems obvious: namely, that choice-worthiness differences are the same between the two views with respect to all moral issues other than the one on which they differ.

The second class is *variable-extension* cases: unlike the former, these are cases involving complete theories, considered in the abstract.<sup>10</sup> Consider, for example, two forms of utilitarianism. They both have exactly the same hedonistic conception of welfare, and they both agree on all situations involving only humans: they agree that one should maximize the sum total of human welfare. They only disagree on the extension of bearers of value. One view places moral weight on animals; the other places no moral weight on animals, and they therefore disagree in situations where animals will be affected. Between these two theories, the intertheoretic comparison seems obvious: they both agree on how to treat humans, and therefore it seems clear that the choice-worthiness difference of saving one human life compared to saving no human lives is the same on both theories. Other similar examples can be given. If we consider a form of utilitarianism that claims that only presently existing people have moral weight and we should maximize the sum of their wellbeing, and compare that to total utilitarianism, again there is an intuitively obvious comparison: the choice-worthiness differences are the same in situations that only affect presently existing people.

and common-sense morality is correct', we should really understand them as saying, 'If animals have rights in the way that humans do, then the difference in choice-worthiness between killing an animal and not-killing an animal (in some particular situation) is much greater than the difference in choice-worthiness between killing an animal and not-killing an animal (in some particular situation) if animals don't have rights and common-sense morality is correct.' Given how laborious this would be to say, it's not surprising that natural English would use a slightly less precise shorthand.

<sup>10</sup> See Tarsney, 'Rationality and Moral Risk', Appendix B and sect. 6.3.1, and Christian Tarsney, 'Intertheoretic Value Comparison: A Modest Proposal', *Journal of Moral Philosophy*, vol. 15, no. 3 (June 2018), pp. 324–44.

The third class of cases are those where we make comparisons of people's moral viewpoints, whether between different people, or within one person's life and across times. Consider, for example, the following statements.

Laura used to think that stealing from big corporations was only mildly wrong, but now she thinks it's outrageous.

James thinks that extramarital sex is a minor wrong, but Jane thinks it's an abomination.

Both of these seem to be perfectly meaningful statements. But they are claims about intertheoretic comparisons. Reflecting on these statements suggests that sometimes we make intertheoretic comparisons by taking a detour via preferences. For example, if the first statement is true, then, assuming Laura is a morally conscientious agent, her preference to not-steal from big corporations has increased in strength, in proportion with her belief about the wrongness of stealing. It's part of common sense that we can make comparisons of preference-strength across people, or across changes in preference within one person's life.<sup>11</sup> But if we can make comparisons of preference-strength when those preferences are in proportion with the agent's moral views, then we can make comparisons of choice-worthiness differences, too.

Our intuitions about intertheoretic comparisons are therefore mixed: in some cases, they seem possible; in other cases, they don't. But this is enough to undermine the arguments for scepticism that were based on appeal to intuitions about particular cases.

The sceptic could respond by trying to debunk the intuitions we've appealed to above. She could argue that, rather than comparisons of choice-worthiness differences, our intuitions are simply tracking the ordinal rank of an option on different theories' choice-worthiness orderings. When we say, 'James thinks that extramarital sex is a minor wrong, but Jane thinks it's an abomination,' we are really saying something like, 'James thinks that extramarital sex is approximately in the 40th percentile of choice-worthiness of options (more severe wrong than jay walking, but not as severe a wrong as lying), whereas Jane thinks it's in approximately the 20th percentile (a more severe wrong than lying, but not as severe a wrong as murder).'

<sup>11</sup> Though of course, it has been the subject of considerable debate within economics. For an overview, see, for example, Ken Binmore, *Rational Decisions*, Princeton, NJ: Princeton University Press, 2009. We are unable to enter into that debate here, so what we say should be taken to be on the assumption that the common-sense view about comparisons of preference-strength is correct.

In order to make true statements such as these, we don't need to make intertheoretic comparisons of choice-worthiness differences.

However, this debunking argument seems to misrepresent our judgments in these cases. In many instances of moral uncertainty, MEC-style reasoning seems plausible. As we argued in Chapter 2, if Sophie is a morally conscientious person and is fairly uncertain about whether animals have rights, then it seems plausible that, by her own lights, she shouldn't order factory-farmed chicken, even if she thinks it's more likely than not that ordering the chicken is slightly more choiceworthy than not ordering the chicken. If we were just talking about ordinal rank when we made what seem like intertheoretic comparisons, however, then MEC would be inapplicable. So the best explanation of the fact that we find MEC-style reasoning plausible is that we can make intertheoretic comparisons.

The reasonable view, then, on the basis of the intuitive evidence, is that sometimes intertheoretic comparisons are obvious; sometimes they are unobvious or perhaps impossible. The *appeal to cases* argument and the *swamping* argument therefore give us no argument for intertheoretic incomparability in general. It may be that, ultimately, we want to reject our intuitions about intertheoretic comparisons as confused. But this is a conclusion that we should only draw at the end of enquiry, after we have tried our best to come up with a general account of intertheoretic comparisons and failed.

Things are different for the *arbitrary unit* argument, which is a perfectly general argument against the possibility of intertheoretic comparisons. In light of the examples given above, however, the *arbitrary unit* argument seems to prove too much. If it were correct, it would show that *no* intertheoretic comparisons are possible. Yet we have seen many cases where they do seem to be possible. So we are left with a puzzle. On the one hand we have an argument that the choice of unit in a theory's choice-worthiness function is arbitrary; on the other hand we have specific cases where the choice of unit seems *not* to be arbitrary. The rest of the chapter will discuss accounts of intertheoretic comparisons that might resolve this puzzle. To this end, let's first consider structural accounts of intertheoretic comparisons.

## II. Structural Accounts

Let us define a structural account of intertheoretic comparisons as follows.

A *structural account* is a way of giving different moral theories a common unit that only invokes structural features of the theories' qualitative

choice-worthiness relation (such as the choice-worthiness of the best option and worst option) or mathematical features of its numerical representation (such as the mean, sum, or spread of choice-worthiness). The identities of particular options have no relevance; only positional properties matter.

If we were to interpret Lockhart's PEMT as an account of how theories actually compare (rather than an account of what to do in conditions when they are incomparable), then it would be a structural account. Similarly, one could go further than we claimed in Chapter 4 and suggest that variance voting is the correct account of how two theories compare, when they are comparable. This would also be a structural account.

Structural accounts are appealing for at least two reasons. First, they confront the 'arbitrary unit' argument for intertheoretic incomparability head on. If some structural account is correct, then we do not require anything more from moral theories other than that they provide an interval-scale measurable choice-worthiness function. In order to normalize two theories, all they need is that those two theories both provide a choice-worthiness ordering over prospects that satisfies the von Neumann–Morgenstern axioms.

Second, they alleviate the swamping worry, too. In virtue of *only* looking at each theory's choice-worthiness function, they attempt to ensure that no theory gets more 'say' than another. Lockhart, for example, explicitly defends the PEMT on these lines:

The PEMT might be thought of as a principle of fair competition among moral theories, analogous to democratic principles that support the equal counting of the votes of all qualified voters in an election regardless of any actual differences in preference intensity among the voters.<sup>12</sup>

We have found that structural accounts have regularly been endorsed in conversation for these reasons. However, though structural accounts have some theoretical appeal, we believe that all such accounts are flawed. To show this, we provide five arguments against structural accounts; moreover, seeing why structural accounts fail will help us to see how non-structural accounts could succeed.

These five objections share a common theme. According to structural accounts, there is only *one* possible way to make intertheoretic comparisons between any two theories. In order to reject structuralism we therefore don't

<sup>12</sup> Ted Lockhart, *Moral Uncertainty and Its Consequences*, p. 86.



need to argue that two theories definitely do compare in a certain way that is inconsistent with any structuralist view. All we need to argue is that an agent is not making a mistake if she has a belief that two theories compare in a way that structural accounts cannot countenance. Each of the five objections that follow show a different way in which this is the case.

### III. Five Arguments against Structural Accounts

#### Varied-Extension Cases

In cases where two theories agree in every respect except on the extension of the class of things that are fundamental bearers of value, there seems to be an intuitive way in which these two theories compare. For example, consider humans-only utilitarianism (HO-U) and all-sentient-creatures utilitarianism (ASC-U). The natural and obvious way to compare these two theories is to suppose that the value of humans is the same according to both theories; certainly it seems epistemically permissible for a decision-maker to have a credence distribution such that this is so. But structural accounts have to deny this. Because structural accounts are only sensitive to features of a theory's choice-worthiness function, they can't 'peer inside' the theories to see that they have parts in common, and they can't allow one theory to be higher-stakes in general than another.

To see this, let's suppose that ASC-U regards (non-human) animal lives as half as valuable as human lives (due to some aspect of their reduced cognitive capacities). And let's suppose, for simplicity, that there are only three possible options: one in which there are no sentient creatures, one in which there are 100 humans, and 0 other animals, and one in which there are 100 humans and 800 other animals. In Table 5.1 those two theories are represented in the most intuitive way.

This fits our intuition that, since they have the same account of human value, the second option is considered just as valuable by each theory. It also fits our intuition that if you lived in the world with many animals and came to change your beliefs from a theory that neglected animals to one that recognized them, you would find the world to be much more valuable than before and to think that there was more at stake.

This, however, is not how structural accounts would normalize them. For example, the broad PEMT would normalize them as in Table 5.2.

**Table 5.1**

	HO-U	ASC-U
0 humans, 0 animals	0	0
100 humans, 0 animals	100	100
100 humans, 800 animals	100	500

**Table 5.2**

	HO-U	ASC-U
0 humans, 0 animals	0	0
100 humans, 0 animals	100	20
100 humans, 800 animals	100	100

The broad PEMT therefore has to deny that the two theories agree on the value of human lives. The same would be true if we normalized the two theories at the difference between the maximum and mean choice-worthiness, at their variance, or at any other features of their choice-worthiness function. Structural accounts have to say that, according to ASC-U, humans are of less value than they are according to HO-U, and that a decision-maker is making a mistake if she believes them to be otherwise. But this seems wrong. They therefore fail to capture our intuitions about intertheoretic comparisons in varied-extension cases.

The structuralist might be inclined to reject the starting intuition we appealed to. But it was intuitions like those concerning varied-extension cases that made us think that intertheoretic comparisons were possible at all. Insofar as structural accounts reject those intuitions, they, therefore, undermine part of the motivation for denying scepticism in the first place.

### Discontinuity with Universal Indifference

A second problem for structural accounts is that it seems, intuitively, that some theories can be higher stakes than others. This means that structural accounts must create a discontinuity between *universal indifference*—the view on which all options are equally choiceworthy—and discerning moral theories.

For example, suppose that Meursault originally adheres to common-sense ethics, but then reads more and more nihilist literature. He becomes progressively convinced that universal indifference is true. However, the way he becomes convinced is not that he increases his credence in universal indifference and decreases his credence in common-sense ethics. Rather, he progressively realizes that certain things he used to think were of positive value are neutral in value. First, he realizes that art and literature and non-personal goods have no intrinsic positive or negative value. Then he realizes that there are no other-regarding reasons, and retreats to egoism. At each step, Meursault becomes more despondent. Finally, he realizes that even his own happiness is also merely neutral in value, and he comes to accept full-blown universal indifference.

A natural and intuitive way to understand this is that the ethical viewpoints that Meursault adheres to become progressively closer and closer to universal indifference. Meursault progressively thinks that there is less and less positive value in the world, until eventually he thinks there is no positive value at all in the world. Again, it seems at least epistemically permissible for Meursault to think about his changes in moral beliefs in this way. However, structuralist accounts cannot understand Meursault's progression in beliefs in this way. According to structuralist accounts, when Meursault rejects the value of art, his other beliefs compensate and he comes to believe that personal moral reasons were much more important than he had previously thought; when Meursault rejects moral reasons, he must also come to believe that his own happiness is much more important than he had previously thought. The amount of value Meursault thinks exists in general is the same right up until the point when he embraces universal indifference. At that point, there is a stark discontinuity. Insofar as structural accounts cannot countenance the possibility that Meursault comes to have progressively lower-stakes beliefs, we have another reason against accepting structural accounts.

### Incoherent Beliefs—Weighing Values

Our third argument against structural accounts rests on the fact that how we think we should compare two theories can be affected by our beliefs in theories other than the two in question.<sup>13</sup>

<sup>13</sup> The idea that pluralistic theories can serve as 'bridges' to establish intertheoretic comparability between other theories, based on cases like the one described below, is defended in Tarsney, 'Rationality and Moral Risk', pp. 202–4 and 323–7.

Let us suppose that Abel has the following beliefs. He is a consequentialist, but isn't sure if pleasurable experiences are of value, or if artistic and scientific accomplishments are of value, or if both are of value, and, therefore, he isn't sure about the relative choice-worthiness of options that promote happiness and those that promote accomplishments. However, he has no doubts at all about how to weigh the value of increased pleasurable experiences against the value of artistic and scientific accomplishments, if they are both of value.

That is, he has credence in three different theories:  $T_1$ , according to which only pleasurable experiences are of value;  $T_2$ , according to which only artistic and scientific accomplishments are of value, and  $T_3$ , according to which both are of value.

Abel believes that the intratheoretic comparison between pleasure and accomplishments on  $T_3$  provides a basis for the intertheoretic comparison between  $T_1$  and  $T_2$ . He believes that one unit of pleasurable experience is worth the same on  $T_1$  and  $T_3$  and he believes that one unit of accomplishment is worth the same on  $T_2$  and  $T_3$ . Does this set of beliefs seem irrational? We believe not. But structural accounts would have to say that these beliefs are irrational.

To illustrate, suppose for simplicity that there are only three possible options, *A*, *B*, and *C*. On  $T_3$ , the pleasure produced by option *A* is of value 100, the achievement produced by option *B* is of value 200, the pleasure produced by option *C* is of value 50, and the achievement produced by option *C* is of value 100. Given our description of Abel's beliefs, we would therefore represent these theories as in Table 5.3.

However, structural accounts will not be able to represent theories in this way. The broad PEMT, for example, will represent them as in Table 5.4.

For any comparison between two theories, structural accounts are blind to the decision-maker's beliefs about other theories. But this seems like a mistake. In the case given above, Abel is sure about how pleasure and accomplishment should be weighed against each other if they are both of

**Table 5.3**

	$T_1$	$T_2$	$T_3$
<i>A</i> : Lots of pleasure, little achievement	100	0	100
<i>B</i> : Lots of achievement, little pleasure	0	200	200
<i>C</i> : Moderate amounts of both pleasure and achievement	50	100	150

**Table 5.4**

	$T_1$	$T_2$	$T_3$
A: Lots of pleasure, little achievement	100	0	0
B: Lots of achievement, little pleasure	0	100	100
C: Moderate amounts of both pleasure and achievement	50	50	50

value. But even though he is certain that if both pleasure and accomplishment are of value they weigh against each other in such-and-such a way, according to structural accounts he must believe that, when it comes to the intertheoretic comparison between pleasure and accomplishment, they *must* weigh against one another in a different way. That seems wrong.

### Incoherent Beliefs—Too Much Comparability

Our fourth argument is that structural accounts generate too much intertheoretic comparability. Structural accounts must claim that *all* moral theories that can be represented by an interval-scale measurable choice-worthiness function are intertheoretically comparable. But that seems much too strong.

Consider Beth, who, like Abel is unsure about whether pleasure or accomplishment or both are of value. However, unlike Abel, she is absolutely certain that if both pleasure and accomplishment are of value, then, because of the nature of those two values, they are absolutely incomparable in value. Like with Abel, we can represent her as having credence in three theories:  $T_1$ , according to which only pleasurable experiences are of value;  $T_2$ , according to which only artistic and scientific accomplishments are of value; and  $T_3$ , according to which both are of value, but the value of each is (intratheoretically) incomparable.

Further, let us suppose that Beth believes that  $T_1$  and  $T_2$  are absolutely incomparable. Given that she believes that, if pleasure and accomplishment are both of value, then they are absolutely incomparable, this seems like a natural and reasonable set of beliefs. But structural accounts cannot make sense of this. For structural accounts,  $T_1$  and  $T_2$  are no different from any other theories that provide choice-worthiness functions, and can be normalized in the same way. But this seems to force Beth to have strange beliefs:

believing that, though pleasure and accomplishment would be absolutely incomparable if they were both of value, nonetheless they are perfectly comparable when considered intertheoretically.

The advocate of structural accounts could respond by supposing that her account is only an account of how values compare across theories when those two theories are intertheoretically comparable; and that the account remains silent on *when* theories are intertheoretically comparable. But this drastically reduces the attractiveness of the structural accounts. One of the virtues of the account was its generality, and the fact that it served as a response to the worry that intertheoretic value comparisons are not possible at all. With this modification, we are left with no way of knowing when, if ever, two theories are intertheoretically comparable. If the structuralist wishes to assuage this worry by developing an additional account of when two theories are intertheoretically incomparable, then they will run into other problems. For that account would have to appeal to information other than information concerning the structure of the theory's choice-worthiness function. Such an additional account would therefore go against the very spirit of structural accounts, and should make us wonder why we were attracted to them in the first place.

### Amplified Theories

Our final argument is that structural accounts can't account for a possible way in which a decision-maker might believe two theories to be related. Introducing some new terminology, let us say that two theories  $T_i$  and  $T_j$  have the same *interval-scale structure* iff there exists a constant  $c$  and a constant  $k > 0$ , such that for all options  $A$ :  $CW_i(A) = k CW_j(A) + c$ , where  $CW_i(A)$  represents the choice-worthiness of  $A$  on moral theory  $T_i$ . And let us say that  $T_i$  is an *amplified* version of  $T_j$  iff they have the same interval-scale structure and the difference in choice-worthiness between any two options  $A$  and  $B$  on  $T_i$  is greater than the difference in choice-worthiness between those same two options  $A$  and  $B$  on  $T_j$ . Figure 5.1 (which is to scale) represents this idea.

Suppose that  $A-E$  are the only possible options. If so, then  $T_2$  is an amplified version of  $T_1$ .  $T_1$  and  $T_2$  agree that the difference between  $B$  and  $C$  is four times the difference between  $A$  and  $B$ . But the difference between  $A$  and  $B$ , according to  $T_2$ , is twice the difference between  $A$  and  $B$ , according to  $T_1$ . So  $T_2$  is an amplification of  $T_1$ .




Figure 5.1

With this on board, we can state our final argument against structuralism.

(P1) It's epistemically permissible to believe in two distinct theories, one of which is an amplified version of the other.

(P2) If (P1), then all structural accounts are false.

(C3) Therefore, all structural accounts are false.

(P2) is uncontroversial. If structural accounts are correct, we can only appeal to information concerning the theory's choice-worthiness function, which is unique only up to a positive affine transformation. So, on structural accounts, all theories with the same interval-scale structure must be normalized in exactly the same way. Providing an example where it seems epistemically permissible to believe that one theory is an amplified version of another theory would thereby show that structural accounts are not correct. Here we suggest an example of such a pair of theories.

### *Sophie's Change of View*

Sophie initially believes in a partialist form of utilitarianism, which posits both impartial and agent-relative value. Though she thinks that human welfare is of value in and of itself, she also thinks that the presence of certain relationships between her and others confers additional value on those with whom she has the relationship. For that reason, she believes that the welfare of her family and friends is more valuable than that of distant strangers, though she thinks that both have value.

Sophie then revises her belief, and comes to believe that the welfare of all humans is of equal value. However, she realizes that there are two ways in which she could come to hold this view. First, she could come to believe that there's no such thing as agent-relative value; no relationships confer additional value on the welfare of others. In which case the value of the welfare of distant strangers would be the same as she had previously thought, but the value of the welfare of close family and friends would be less than she had previously thought. Second, she could come to believe that, morally, she should 'be a brother to all', and she should regard her relationship with all other humans as being morally valuable in just the same way that she had thought that blood relationships and friendships were morally valuable. In which case, the welfare of her family and friends would be just as valuable as she had always thought; it's just that the value of the welfare of distant strangers is greater than she had thought. She is unsure which she should believe.

Let's call the first view that Sophie considers *Benthamite utilitarianism* and the second view *kinship utilitarianism*. Intuitively, it seems perfectly meaningful to think that Sophie could be uncertain between these two views. And it also seems meaningful for her to think that her relationships would have been downgraded in value, if Benthamite utilitarianism were true, but that the value of distant strangers would have increased in value, if kinship utilitarianism were true.

One might think that there is no *meaning* to the idea of one theory being an amplified version of another theory. But we can point to five distinctions between the two theories in order to further explain the meaningfulness of amplified theories.

First, Benthamite utilitarianism and kinship utilitarianism differ on the grounding of choice-worthiness: they disagree on facts concerning in virtue of what are certain actions wrong. Benthamite utilitarianism would claim that saving a human life is good because saving that life would increase the sum total of human welfare. On Benthamite utilitarianism, there is just one fact in virtue of which saving a human life is a good thing. In contrast, kinship utilitarianism would claim that saving a human life is good both because saving that life would increase the sum total of human welfare, and also because one has a certain sort of relationship to that person. On kinship utilitarianism, there are two facts in virtue of which saving a human life is a good thing. That is, Benthamite utilitarianism and kinship utilitarianism disagree on what the *right-makers* are.



In general, there is often more to a moral theory than a choice-worthiness function: there is also a metaphysical account of why that choice-worthiness function is correct. This provides some grounds for thinking that the *arbitrary unit* argument, which appealed to the idea that there is nothing more to a moral theory than its choice-worthiness function, is mistaken. Theories differ in their metaphysics, and, intuitively, that metaphysical account can make a difference to the amplification of a theory. On Benthamite utilitarianism, one does *one* wrong thing by killing another person (namely, reducing the amount of welfare in the world), whereas, on kinship utilitarianism, one does *two* wrong things (reducing the amount of welfare in the world, and violating an obligation that arises out of a special relationship that one has). Committing both wrong *X* and wrong *Y* is worse than committing just wrong *X*. So it's a more severe wrong to kill, according to kinship utilitarianism, than it is according to Benthamite utilitarianism.

A second way in which we can make sense of amplified theories is with reference to the relationships in which they stand to other theories. Benthamite utilitarianism and kinship utilitarianism differ in their relationship to partialist utilitarianism. Benthamite utilitarianism has a part in common with partialist utilitarianism—the part that concerns strangers. Kinship utilitarianism also has a part in common with partialist utilitarianism, but it is a different part: the part that concerns family and friends. Because of these different relationships, we can make sense of kinship utilitarianism being an amplified version of Benthamite utilitarianism.

Third, it seems that which attitudes it is fitting for Sophie to have, given revision of her initial belief, depends on which amplification of utilitarianism she comes to believe. If she comes to believe Benthamite utilitarianism, it seems fitting for her to be disappointed: she has lost something of value, as her family and friends are merely as valuable as distant strangers. In contrast, the same is not true if she comes to believe kinship utilitarianism. Perhaps, instead, it would be fitting for her to feel a sense of wonder and new connectedness with those whom she doesn't know.<sup>14</sup>

Fourth, it seems plausible to us that the epistemological facts can differ depending on which theory we are discussing, and that they can differ in virtue of the amplification of the theory. Perhaps the idea of downgrading

<sup>14</sup> Note that we use the term 'fitting' rather than 'ought'. That an attitude is fitting does not entail, necessarily, that one ought to have that attitude. Analogously, one might reject the idea that the requirements of etiquette affect what you ought to do while still acknowledging that it's against the requirements of etiquette to eat with one's elbows on the table.

the value of her family and friends seems abhorrent to her; or perhaps she finds the idea that certain relationships should confer additional value on welfare metaphysically spooky. Either of those views seem reasonable, and either one would mean that she'd find one of the two theories more plausible than the other.

Fifth, facts about what it's appropriate to do under moral uncertainty can differ depending on which amplification of utilitarianism Sophie has credence in. If she has 20% credence in kinship utilitarianism and 80% credence in non-consequentialism, then, if she follows MEC, she will more often act in accordance with utilitarianism than if she has 20% credence in Benthamite utilitarianism and 80% credence in non-consequentialism. This is because things are higher-stakes in general for kinship utilitarianism than for Benthamite utilitarianism.

One might complain that we have only given one example, and that we shouldn't trust our intuitions if they pertain to merely one case. But we could give more examples. Consider Thomas, who initially believes that human welfare is ten times as valuable as animal welfare, because humans have rationality and sentience, whereas animals merely have sentience. He revises this view, and comes to believe that human welfare is as valuable as animal welfare. He might now think that human welfare is less valuable than he previously thought because he has rejected the idea that rationality confers additional value on welfare. Or he might now think that animal welfare is more valuable than he previously thought, because he has extended his concept of rationality, and thinks that animals are rational in the morally relevant sense.

Or consider Ursula, who initially believes that wrong acts are ten times as wrong as wrong omissions, but then comes to believe that acts and omissions are on a par. Does she come to believe that wrong omissions are worse than she had thought, or does she come to believe that wrong acts aren't as wrong as she had thought? If the former, then it might be fitting for her to feel horror at the idea that, insofar as she had let others die, she had been doing things as bad as murder all her life. If the latter, then it might be fitting for her to feel less blame towards those who had killed others.

In exactly the same way as with Sophie, we can explain the distinction between these pairs of amplified theories by looking at differences in rightmakers, differences in fitting attitudes, differences in epistemological reasons, and differences in facts about what it is appropriate to do under moral uncertainty.

For these reasons, we believe that we should reject structural accounts of intertheoretic comparisons.

#### IV. Non-structural Accounts

In the course of our arguments in favour of amplified theories, we saw that two theories with the same interval-scale structure can differ in a number of ways—their metaphysical underpinnings, their relationship to other theories, their relationship to epistemic reasons, the reactive attitudes that are fitting, and the actions that are rational given credence in them—and that these differences have some relationship to intertheoretic comparisons. Each of these has the potential to enter into an explanation of the possibility of intertheoretic comparisons. We could say, for example, that such-and-such an intertheoretic comparison is true *because* of certain facts about what attitudes it is fitting to have;<sup>15</sup> or we could say that it's true *because* of facts about what it's rational to do under moral uncertainty.<sup>16</sup> And in the course of some of our other arguments, we saw that sometimes we can make intertheoretic comparisons via comparisons of preference-strength. If we want to compare  $T_1$  and  $T_2$ , perhaps we can do so by comparing the preference-strengths of morally conscientious person *A*, who fully believes  $T_1$ , and morally conscientious person *B*, who fully believes  $T_2$ . And we sometimes saw that we can appeal to relationships between theories—if we can explicate the notion of some aspect of a theory being 'shared' across two theories, then again we would have a way of making intertheoretic comparisons.<sup>17</sup>

Our problem, therefore, is not that we have no way of making the comparison, but that we have too many. There are many ways in which theories differ that seem to relate to intertheoretic comparisons. But we don't yet know which of these aspects are the grounds of intertheoretic comparisons, and which are consequences of intertheoretic comparisons.

Even without having a specific account in hand, however, we may have the basis for optimism about the extent of the applicability of MEC. If we allow the possibility of amplified theories, then we should reconsider what

<sup>15</sup> This account is suggested by Sepielli, "Along an Imperfectly Lighted Path".

<sup>16</sup> This account is suggested by Ross, 'Rejecting Ethical Deflationism'.

<sup>17</sup> This idea is suggested by Ross, 'Rejecting Ethical Deflationism' and then explicated by Sepielli, 'What to Do When You Don't Know What to Do' (who recants the view in "Along an Imperfectly Lighted Path").

we call moral theories. Rather than thinking of ‘utilitarianism’ as designating one particular theory, really it designates an entire class of theories, each of different levels of amplification. We can therefore return to the ‘hard cases’ for intertheoretic comparisons with a new perspective. For example, rather than thinking that there is simply no way to make the comparison between utilitarianism and prioritarianism, we might instead think that we are just unsure about which, of all the prioritarian theories within the class of theories with the same interval-scale structure, is most plausible. Some forms of prioritarianism clearly seem implausible, such as the form of prioritarianism according to which the value of one extra year of healthy life given to a typical member of an affluent country is one million times as large as the value of one extra year of healthy life given to a typical member of an affluent country according to utilitarianism. When we were initially thinking about the comparison between utilitarianism and prioritarianism, the argument was that, because there was no privileged way to make the comparison, we should conclude that there is no comparison. But in light of the discussion of amplified theories, the lack of a privileged normalization shouldn’t be so worrying to us. Instead, we should distribute our credences over many different prioritarianisms with the same interval-scale structure. And we have intuitions about that: it’s clear we should have much higher credence in the prioritarianism that values one year of healthy life given to a typical member of an affluent country approximately the same as utilitarianism does than we should to the prioritarianism that values that year of life as one million times as much as utilitarianism does. But if we even have a probability distribution over different prioritarianisms of different levels of amplification, that’s sufficient to use MEC.

However, the arguments that we’ve given might make us even more worried than we were by swamping. Consider, for example, someone who is unsure between prior-existence utilitarianism and total utilitarianism. If our arguments are correct, the natural way to normalize these two theories is via the part on which they agree, namely the value of presently existing people. However, if so, then it seems that total utilitarianism will swamp prior-existence utilitarianism: whereas the number of people who presently exist is 7 billion, it seems that the expected number of people who exist in the future is at least in the tens of trillions.<sup>18</sup> For almost any decision that

<sup>18</sup> For context, *Homo sapiens* have already been around for 200,000 years and the average mammalian species lasts for 1 to 2 million years. If we had even a one in ten chance of surviving for as long as a typical mammalian species, that would be an expected 10 to 20 trillion

has some effect on the long-term future of the human species, the action with the highest expected choice-worthiness will accord with total utilitarianism rather than present-existence utilitarianism, even if one has a very small credence in total utilitarianism.

However, we think that the correct response is to deny that the *swamping* argument is a good argument against certain intertheoretic comparison claims. If we take seriously the idea that there are norms governing decision-making under moral uncertainty, then presumably we do so because of the analogies between decision-making under moral uncertainty and decision-making under empirical uncertainty. But a *swamping* argument against an empirical hypothesis would be absurd. For example, prior to the first atomic test, physicist Edward Teller raised the possibility that a nuclear blast might 'ignite' the atmosphere by causing a self-propagating fusion reaction between nitrogen nuclei, thereby destroying the planet.<sup>19</sup> Prior to extensive physical calculations, this hypothesis would have swamped the expected utility calculation of undertaking a first atomic test. But that wasn't a reason for supposing that the destruction of the planet wouldn't be as bad, if Teller's hypothesis is true, than if Teller's hypothesis is false. If one empirical hypothesis regards a decision-situation as higher-stakes than another, we should represent it as such; the same is true for moral theories. A theory on which there is more to gain or lose in a situation (or in general) should have greater weight in the calculation concerning what to do: that's precisely the point of taking into account both the probability that the theory is true *and* the strength of the choice-worthiness differences according to the theory.

Partly, we think that our concerns about swamping are influenced by concerns about whether maximizing an expectation is the right way to make decisions in the face of tiny-probability but extremely-high-value outcomes; this is the 'fanaticism' problem that we discuss in Chapter 6. We agree that it is a worrying problem that maximizing an expectation might require one to pursue actions that have almost no chance of producing a good outcome. But this is a problem for decision theory in general, whether under moral or

more individuals. Given any real chance of spreading beyond Earth, the expected number would be much higher.

<sup>19</sup> See Richard Rhodes, *The Making of the Atomic Bomb*, London: Simon & Schuster, 1986, pp. 418–19.

empirical uncertainty.<sup>20</sup> It has nothing to do with intertheoretic comparisons in particular.

Given this, it seems that the class of non-structural accounts of intertheoretic comparisons is promising to explore. We can divide non-structural accounts into two categories. According to what we call *common ground* accounts, intertheoretic comparisons are true in virtue of different theories' having parts that are shared between them. Ross, Sepielli, and Tarsney have proposed common ground accounts.<sup>21</sup>

According to *universal scale* accounts, intertheoretic comparisons are true in virtue of the fact that there is some independent choice-worthiness scale that is the same across different theories. Ross and Sepielli have both proposed universal scale accounts, in addition to the common ground accounts that they have proposed. We defend a different universal scale account later in this chapter.

We can represent these different accounts diagrammatically. Let us consider two theories,  $T_1$  and  $T_2$  (see Figure 5.2).

Structural accounts normalize with respect to some features of each theory's choice-worthiness function. In the diagram below, we have normalized them with respect to the range of the choice-worthiness function. The key question for structural accounts is at which features of each theory's




Figure 5.2

<sup>20</sup> See, for example, Alan Hájek, 'Waging War on Pascal's Wager', *The Philosophical Review*, vol. 112, no. 1 (January 2003), pp. 27–56 and Nick Bostrom, 'Pascal's Mugging', *Analysis*, vol. 69, no. 3 (July 2009), pp. 443–5.

<sup>21</sup> Ross, 'Rejecting Ethical Deflationism', pp. 764–5; Sepielli, 'What to Do When You Don't Know What to Do'.




Figure 5.3

choice-worthiness function to normalize (such as the range, or the variance, or the maximum choice-worthiness minus the mean choice-worthiness).

Common ground accounts attempt to find some choice-worthiness-differences between specific options that are agreed on by both theories. As opposed to structural accounts, common ground accounts require us to be able to identify options across theories (rather than merely identifying them by their position in the choice-worthiness function). The key questions for common ground accounts are: (i) to elucidate what it means for a theory to ‘share parts’; and (ii) to identify the options *A* and *B* whose choice-worthiness difference the two theories under consideration agree upon. In Figure 5.3, we have supposed that the two theories agree on the choice-worthiness difference between *A* and *B*.

According to universal scale accounts, the two theories are already plotted on some shared scale, represented in black in Figure 5.4.

The key question for universal scale accounts is to explain the nature of this shared scale, and give reasons for thinking that this shared scale exists.

## V. Against Two Common Ground Accounts

On a common ground account that is suggested by both Ross and Sepielli,<sup>22</sup> the idea, in Ross’s words, is to look at ‘cases in which, for some pair of

<sup>22</sup> Ross, ‘Rejecting Ethical Deflationism’, pp. 764–5; Sepielli, ‘What to Do When You Don’t Know What to Do’.




Figure 5.4

options, we know that the difference between their values is the same according to both ethical theories.’<sup>23</sup> We then can use that difference to define one unit of choice-worthiness that is comparable across both theories.

The trouble with this account is that neither Ross nor Sepielli give an explanation of what it is for some choice-worthiness difference to be ‘shared’ between two options. Sepielli is clearest: he takes agreement between theories to consist in the fact that two theories agree where some part of their choice-worthiness functions have the same interval-scale structure. More precisely, Sepielli’s view is as follows. For some three particular options *A*, *B*, and *C*:

$$\text{If } \frac{CW_i(A) - CW_i(B)}{CW_i(B) - CW_i(C)} = \frac{CW_j(A) - CW_j(B)}{CW_j(B) - CW_j(C)}$$

$$\text{then } CW_i(A) - CW_i(B) = CW_j(A) - CW_j(B).$$

But this account is internally inconsistent, and therefore the claim above is false.<sup>24</sup> We saw this in Chapter 4, section I, with respect to utilitarianism

<sup>23</sup> Ross, ‘Rejecting Ethical Deflationism’, p. 764.

<sup>24</sup> Sepielli recants this view, because of this objection, in Sepielli, ‘‘Along an Imperfectly Lighted Path’’.



and prioritarianism (where we assumed the prioritarian's concave function is the square root function). To recap: suppose that Annie and Betty have lived for sixteen years so far, and if you save their lives they'll each live a further nine years. Both utilitarianism and prioritarianism agree that the difference in choice-worthiness between saving both Annie and Betty and saving Annie only is the same as the difference in choice-worthiness between saving Annie only and saving neither. According to the prioritarian, the choice-worthiness difference between saving Annie only and saving neither is  $\sqrt{25} - \sqrt{16}$ , which equals 1. According to the utilitarian, the difference is  $25 - 16$ , which equals 9. So, according to Ross's and Sepielli's view, 1 unit of choice-worthiness on prioritarianism equals 9 units of choice-worthiness on utilitarianism.

But now suppose that Annie and Betty had lived for sixty-four years, and would live a further nine years. Again, both utilitarianism and prioritarianism agree that the difference in choice-worthiness between saving both Annie and Betty and saving Annie only is the same as the difference in choice-worthiness between saving Annie only and saving neither. But, according to the prioritarian, the choice-worthiness difference between saving Annie only and saving neither is  $\sqrt{73} - \sqrt{64}$ , which is approximately 0.5. According to the utilitarian, the difference is 9, as before. So, according to Ross's and Sepielli's view, 1 unit of choice-worthiness on prioritarianism equals 18 units of choice-worthiness on utilitarianism. But this is different from what we concluded in the previous paragraph, when Annie and Betty had lived shorter lives. So Ross's and Sepielli's account generates inconsistent pronouncements about how choice-worthiness compares across two theories. So their account should be rejected.

Christian Tarsney suggests a variant of Ross and Sepielli's account.<sup>25</sup> On Tarsney's view, the common ground between different theories is not some specific ratio of choice-worthiness differences, but instead is a shared *category of reasons*. So, for example, if we consider one theory that values both pleasure and beauty, and another theory that values both beauty and knowledge, then *ceteris paribus*, we should think that these two theories agree on how strong the reasons to promote beauty are.

One might think that this suffers from internal inconsistency in much the same way that Ross and Sepielli's accounts do. Consider a decision-maker who has 1/3 credence in each of the following three theories:  $T_1$  values pleasure and knowledge;  $T_2$  values knowledge and beauty;  $T_3$  values

<sup>25</sup> Tarsney, 'Rationality and Moral Risk', Appendix B.

beauty and pleasure. Let's call the relevant choice-worthiness units hedons, epistemons, and aesthetons. And let's suppose that, on  $T_1$ , increasing the value of the world by 1 hedon (and leaving everything else as it is) is as choiceworthy as increasing the value of the world by 1 epistemon (and leaving everything else as it is). Similarly, the exchange rate on  $T_2$  between epistemon and aesthetons is 1:1, but on  $T_3$  the exchange rate is 1 aestheton to 2 hedons.

There seems to be nothing irrational with having such a credence distribution. But if a decision-maker did have such a credence distribution, then it would seem that Tarsney's 'categories of reasons' would give inconsistent conclusions on how  $T_1$ ,  $T_2$  and  $T_3$  compare with each other.

However, the precise account that Tarsney defends does not suffer from this problem, because he claims that (i) a decision-maker should only *ceteris paribus* believe that the value of a hedon is independent of what other bearers of value there are, but that (ii) if there is a difference in how hedons and other bearers of value trade off against each other, the decision-maker needs to have some belief that explains why hedons have more or less value on one theory than another. So, on Tarsney's account, insofar as the above decision-maker lacks such a belief, she is indeed irrational.

However, it seems to us that this is now just begging the question. In order to make his account consistent, Tarsney has required the decision-maker to have beliefs about how theories intertheoretically compare (namely, that, *ceteris paribus*, one hedon is worth the same across all theories). Whereas the very claim of the intertheoretic comparability sceptic is that a decision-maker needn't have any such belief.

Indeed, we believe that Tarsney misrepresents the dialectic on the problem of intertheoretic comparisons. Tarsney argues that, 'Opponents of any kind of intertheoretic comparability... must hold that Alice should accept [that the value of a hedon if aesthetons have non-derivative value is incomparable with the value of a hedon if aesthetons don't have non-derivative value].' But in our view, the aim of an account of intertheoretic comparability is not to merely show that the positive position that 'theories are intertheoretically incomparable' is as justified or unjustified as any other position. The aim is to provide an account of *why* intertheoretic comparisons hold, if they do. (In the same way that the aim of responding to the external-world sceptic is to help us understand how we know we have hands, not to simply show that believing we're a brain in a vat is also an unjustified position.) Claiming that decision-makers do (or ought to) have certain beliefs in intertheoretic comparisons does not help us in

that end. So we ultimately find this account, at least as Tarsney has currently developed it, unsatisfactory.

Another common ground account that Sepielli briefly suggests<sup>26</sup> is that there might be ‘paradigms’ of morally acceptable actions, and paradigms of morally heinous actions, which are definitive of choice-worthiness. So just as one might think that the International Prototype Kilogram defines what it means to have 1 kg of mass, so one might think that the difference in choice-worthiness between some two particular, extremely well-specified options (listening to music, and killing for fun, for example), defines one unit of choice-worthiness.

The problem with this account, as a fundamental explanation of how intertheoretic comparisons are possible, is just that there is far too much disagreement among moral theories for this to be a plausible general view. According to ethical egoism, the difference in choice-worthiness between listening to music and killing for fun will be very different compared to the difference in choice-worthiness between listening to music and killing for fun, according to utilitarianism.

The same will be true for any pair of options. One might think that the difference in choice-worthiness between options that only affect the agent—such as, for me, the option to drink a cup of tea right now and the option to keep writing—should be considered the same across all possible moral theories. But this account will result in a clash between separable and non-separable moral views: for example, on average utilitarianism the difference in value between my drinking an enjoyable cup of tea and continuing writing is smaller the larger the number of people there are (because that action affects the average wellbeing by a smaller amount the more people there are); but on total utilitarianism, the value remains constant. So we would have to pick one particular population size in order to put average and total utilitarianism on the same scale, and it’s hard to see how this could be done in a non-arbitrary way.

## VI. Against Two Universal Scale Accounts

The discussion of amplified theories made some suggestions about ways in which we can tell the difference between two theories with the same

<sup>26</sup> Sepielli, ‘“Along an Imperfectly Lighted Path”’, p. 186.

interval-scale structure. This idea motivates some different accounts of intertheoretic comparisons.

We mentioned that the amplification of a theory can make a difference to facts concerning what it's appropriate to do under moral uncertainty. So perhaps it's those facts that make it the case that a certain intertheoretic comparison holds. This is a view suggested by Ross.<sup>27</sup>

As we understand this suggestion, the claim is that facts about how choice-worthiness differences compare across theories are determined by facts about what it is appropriate to do in light of uncertainty between those theories. If an agent faces options  $A$  and  $B$ , and has 10% credence in  $T_1$ , according to which  $CW_1(A) > CW_1(B)$ , and 90% credence in  $T_2$ , according to which  $CW_2(B) > CW_2(A)$ , and it is appropriate for her to do  $A$ , then, *because* it is appropriate for her to do  $A$  in this situation,  $(CW_1(A) - CW_1(B))$  is at least 9 times greater than  $(CW_2(B) - CW_2(A))$ .

The obvious objection to this account is that it puts the cart before the horse. Consider Kate, who has 80% credence in common-sense views about how she should spend her money, and 20% credence in Singer's view that she has strong obligations to donate much of her money to alleviate extreme poverty. In this case, intuitively it's appropriate for her to donate the money. But we have that intuition because it seems clear how the choice-worthiness differences compare across the two moral views in which she has credence. It's not that we have the intuition that it's appropriate for Kate to donate part of her income, and thereby infer what the respective choice-worthiness differences between the common-sense view and Singer's view are. Ross's proposal therefore seems to get the order of explanation the wrong way around.

A different sort of meta-scale account is suggested by Sepielli.<sup>28</sup> He wishes to use degrees of blameworthiness as the scale by which choice-worthiness difference may be compared. The exact nature of his proposal is unclear. But it seems to us that his principal initial proposal is that a decision-maker believes that  $(CW_i(A) - CW_i(B)) = (CW_j(C) - CW_j(D))$  iff the strength of the decision-maker's disposition to blame for doing  $A$  rather than  $B$ , conditional on  $T_i$ , is the same as the strength of the decision-maker's disposition to blame for doing  $C$  rather than  $D$ , conditional on  $T_j$ . It should be fairly clear that this isn't the right account. The decision-maker might just have the sort

<sup>27</sup> Ross, 'Rejecting Ethical Deflationism', p. 763. It is also endorsed by Stefan Riedener ('Maximising Expected Value under Axiological Uncertainty') and by John Broome (private communication, June 2013).

<sup>28</sup> Sepielli, "Along an Imperfectly Lighted Path", p. 183.

of personality where she wouldn't be terribly disposed to blame, if some very demanding moral theories were true. Or it might be that she would be deeply depressed if one particular theory were true, and therefore her dispositions to do anything would be weaker than they ordinarily are. But these factors don't seem to affect how choice-worthiness differences compare across the different theories in which she has credence.

One might try to tweak Sepielli's account by claiming that choice-worthiness differences are measured by how disposed to blame one *ought* to be. But that account would suffer from problems as well. On utilitarianism, how disposed to blame one *ought* to be is not perfectly correlated (indeed, is sometimes highly uncorrelated) with the degree of wrongness of a particular action. So this account would misrepresent choice-worthiness differences according to utilitarianism.

Instead, the best account in this area, we think, is that choice-worthiness differences are measured by the degree to which it is *fitting* to blame for a certain action (or, as we will use the term, the degree to which an action is *blameworthy*). More precisely:  $(CW_i(A) - CW_i(B)) = (CW_j(C) - CW_j(D))$  iff the blameworthiness of the decision-maker for doing *B* rather than *A*, conditional on  $T_i$ , is the same as the blameworthiness of the decision-maker for doing *D* rather than *C*, conditional on  $T_j$ . Note that, on this view, fittingness is a *metaethical* fact: on the assumption of a certain metaethics being true, the relationship between fittingness-to-blame and strengths of reasons is not something that different moral views can disagree about.

We think that this account has at least something going for it: in our discussion of amplified theories, we suggested that there is a link between the amplification of a theory and which attitudes it is fitting to have. The principal question, again, however, is whether choice-worthiness differences should be explained in terms of fitting attitudes, or the other way around. And this fitting-attitude account suffers from the following problem, which is that it cannot explain where interval-scale measurable degrees of blameworthiness come from.<sup>29</sup> It cannot, for example, use probabilities to provide the interval-scale measure. To do so would require making claims such as:

*S* is equally blameworthy for choosing (i) *A* and the guarantee that  $T_1$  is true, as she is for choosing (ii) a 50% probability of *B* and  $T_2$  being true, and a 50% probability of *C* and  $T_2$  being true

<sup>29</sup> We owe this point to Riedener, 'Maximizing Expected Value under Axiological Uncertainty'.

(which would show that the difference in choice-worthiness between  $T_2(B)$  and  $T_1(A)$  is the same as the difference between  $T_1(A)$  and  $T_2(C)$ ). But if ‘probability’ in that sentence means objective chance, then it doesn’t make any sense, because there can’t be objective chances about which theories are true (except 1 and 0). If ‘probability’ means either ‘subjective credence’ or ‘rational credence’, then the account becomes extremely similar to Ross’s ‘facts about appropriateness’ account, which, as we saw, got the order of explanation the wrong way around. So we don’t think that this account is satisfactory, either.

## VII. A Universal Scale Account

We believe that we can make progress on understanding intertheoretic comparisons by learning from work that has been done in the literature on the metaphysics of quantity. Indeed, we seem to be able ask very similar questions about intertheoretical comparisons of quantities and intertheoretical comparisons of choice-worthiness.

P1. Does physical theory  $P_1$  assign greater mass to object  $x$  than  $P_2$  does?

P2. Does physical theory  $P_1$  assign twice as much mass to object  $x$  than  $P_2$  does?

P3. Does  $P_1$  assign a greater difference in mass to  $x$  and  $y$  than  $P_2$  does?

T1. Does moral theory  $T_1$  assign greater choice-worthiness to option  $A$  than  $T_2$  does?

T2. Does moral theory  $T_1$  assign twice as much choice-worthiness to option  $A$  as  $T_2$  does?

T3. Does moral theory  $T_1$  assign a greater difference in choice-worthiness to  $A$  and  $B$  than  $T_2$  does?

If the questions about intertheoretical comparisons of mass (P1–P3) are meaningful, which they surely seem to be, why should we not say the same thing about the questions about the intertheoretical comparison of choice-worthiness (T1–T3)?

The debate around the metaphysics of quantity addresses questions such as: ‘In virtue of what is this object more massive than this other object?’ or ‘In virtue of what is it true that this object is 2 kg and that object is 4 kg?’ There are two classes of answers. Comparativists answer that it is the

mass-relations (' $x$  is more massive than  $y$ ') that are fundamental, and claims about intrinsic mass properties (' $x$  is 4 kg') are grounded in mass-relations. Absolutists answer that it is the intrinsic mass properties of objects that ground mass-relations. For absolutists, the fact that  $x$  is heavier than  $y$  is true in virtue of facts about the intrinsic properties of the objects themselves; for comparativists, it is the other way around.

Though work on the metaphysics of quantity has, so far, entirely focused on scientific quantities ('mass', 'size', 'temperature', etc), we can ask just the same questions about the metaphysics of quantities of value, or of choice-worthiness. We can ask: If it is true that the difference in choice-worthiness between  $A$  and  $B$  is twice as great as the difference in choice-worthiness between  $B$  and  $C$ , is that true in virtue of the fact that  $A$ ,  $B$  and  $C$  each have an intrinsic property of a certain degree of choice-worthiness? Or is the metaphysical explanation the other way around? Moreover, in the same way as the possibility of amplified theories is a crucial issue in the debate concerning intertheoretic comparisons, the possibility of a world in which everything is identical except insofar as everything is twice as massive is a crucial issue in the debate between absolutists and comparativists.<sup>30</sup>

Within the metaphysics of quantities literature, it is generally recognized that absolutism is the more intuitive position.<sup>31</sup> Yet it seems to us that all the discussion of intertheoretic comparisons so far has assumed comparativism about quantities of value or choice-worthiness. If we reject that assumption, then we can provide a compelling metaphysical account of intertheoretic comparisons. In what follows, we'll first present the comparativist account of mass, then quickly sketch Mundy's elegant absolutist account of mass, then explain how something like this account could be applied to value and choice-worthiness.

The standard comparativist account of mass is to analyze mass in terms of the relation ' $x$  is more massive than  $y$ ', and the concatenation operator ' $x$  and  $y$  are together equally as massive as  $z$ '. Three things are important to note about standard comparativist accounts. First, the account is first-order: the variables,  $x$ ,  $y$ , and  $z$  are variables over *objects* (rather than over properties, which would make the account second-order). Second, for this reason, the account is nominalist: it gives an account of mass without any reference to the properties of objects. And, third, the account is empiricist: attempting

<sup>30</sup> Shamik Dasgupta, 'Absolutism vs Comparativism about Quantity', *Oxford Studies in Metaphysics*, vol. 8 (2013), pp. 105–48.

<sup>31</sup> See, for example, Shamik Dasgupta, 'Absolutism vs Comparativism about Quantity'.

to give an analysis of mass solely in terms of observable mass-relations. (So, for example, both 'x is more massive than y' and 'x and y are equally as massive as z' can be defined operationally, identifying them with the behavior of those objects on scales: x is more massive than y iff, when x and y are placed on opposite sides of the scale, the scale will tip in x's direction; x and y are together equally as massive as z iff, when x and y are placed on one side of the scale, and z on the other side, then the scale will not tip in either direction.) Using those two relations, and several axioms,<sup>32</sup> it can be shown that the 'x is more massive than y' relation can be represented using numbers, where  $M(x) > M(y)$  iff x is more massive than y, where the numerical representation is unique up to a similarity transformation ( $f(x) = kx$ ).<sup>33</sup>

In contrast, Mundy's<sup>34</sup> account is second-order, defined over properties as well as objects. Letting X refer to the mass of x and Y refer to the mass of y (etc.), the fundamental mass relations, on Mundy's account, are 'X is greater than Y' and 'X and Y are equal to Z'. That is, the fundamental mass-relations are defined over the mass-properties of objects, rather than over those objects themselves. It is therefore clearly realist rather than nominalist: it posits the existence of properties (which are abstract entities), over and above the existence of objects. And it is Platonist rather than empiricist, because properties are abstract entities that can exist without being instantiated. Using this framework, Mundy is able to give a full formal account of quantities of mass; he then argues that there are significant *empirical* reasons for preferring it to the traditional, first-order, comparativist accounts. In

<sup>32</sup> We will use the axiomatization given in Patrick Suppes and Joseph Zinnes, 'Basic Measurement Theory', in R. Duncan Luce, Robert R. Bush, and Eugene Galanter (eds), *Handbook of Mathematical Psychology*, New York: John Wiley & Sons, 1963, vol. 1, pp. 3–76. Let A be the set of all objects, and let 'Rxy' mean 'x is either less massive or equally as massive as y'. Let  $x \circ y$  refer to a binary operation from  $A \times A$  to A: the 'concatenation' of x and y (where 'concatenation' of x and y may be defined as, for example, placing x and y on the same side of a scale). The axioms are as follows.

1. Transitivity: If Rxy and Ryz, then Rxz.
2. Associativity:  $(x \circ y) \circ z = x \circ (y \circ z)$ .
3. Monotonicity: If Rxy then  $R(x \circ z)(z \circ y)$ .
4. Restricted Solvability: If not Rxy, then there is a z such that  $Rx(y \circ z)$  and  $R(y \circ z)x$ .
5. Positivity: Not  $x \circ y \circ Rx$ .
6. Archimidean: If Rxy, then there is a number n such that  $Ry(nx)$  where the notation (nx) is defined recursively as follows:  $1x = x$  and  $nx = (n - 1)x \circ x$ .

As Suppes and Zines note, axiom 5 in conjunction with the order properties of R and the definition of  $\circ$  imply that the set A is infinite.

<sup>33</sup> See David H. Krantz et al., *Foundations of Measurement*, New York: Academic Press, 1971, vol. 1.

<sup>34</sup> For Mundy's full account, see Brent Mundy, 'The Metaphysics of Quantity', *Philosophical Studies*, vol. 51, no. 1 (1987), pp. 29–54.



particular: in order to prove the representation theorem that mass is representable on a ratio-scale, the traditional comparativist account of mass needs to assume that, for any two objects, there is an actual third object that is equal in mass to those two objects. But the universe may well be finite, and if so then this assumption would be false. But it seems very plausible that objects have mass-quantities *whether or not* the universe is finite.

There is considerable debate between absolutists and comparativists. The key issue, however, when it comes to quantities of value or of choice-worthiness, is that absolutism about quantities of choice-worthiness can neatly solve the problem of intertheoretic choice-worthiness comparisons. And, going further, we can develop an analogue of Mundy's account and solve the intertheoretic comparisons problem by appeal to a *second order* universal scale, which measures primitive abstract degrees of choice-worthiness.

Consider the issue of whether there could be a world  $w_1$ , where all the relations between objects are the same as in world  $w_2$ , but where all objects are twice as massive in  $w_1$  as they are in  $w_2$ . It is generally regarded as a problem for comparativism that it cannot make sense of the idea that  $w_1$  and  $w_2$  could be distinct worlds: the mass-relations between all objects in  $w_1$  are the same as the mass-relations in  $w_2$ , so, according to comparativism, there is no difference between those two worlds. In contrast, absolutism is able to explain how those two worlds are distinct. Properties necessarily exist; so the two worlds differ in the intrinsic properties that objects in those two worlds instantiate. Note, also, that, if  $w_1$  and  $w_2$  are distinct worlds, then we have conclusive evidence for the existence of inter-world mass relations: we can say that object  $x$  in  $w_1$  is twice as massive as it is in  $w_2$ .

Similarly, now, consider the issue of whether there could be two theories  $T_1$  and  $T_2$ , where  $T_1$  has the same interval-scale structure as  $T_2$ , but where the choice-worthiness differences between all options are twice as great on  $T_1$  as they are on  $T_2$ . In our argument against structural accounts of intertheoretic comparisons, we argued that this is a genuine possibility. But, if so, then we have a good argument against comparativism about choice-worthiness, according to which the only fundamental facts about choice-worthiness are facts about choice-worthiness relations between options. (One could try to explicate this idea in comparativist terms using Ross's universal scale account; but we saw that that account was unsatisfactory, getting the order of explanation the wrong way around.) In contrast, if we endorse absolutism about choice-worthiness, then we have an explanation for how  $T_1$  and  $T_2$  could be distinct theories. The same choice-worthiness quantities exist in many different epistemically possible worlds, so we can

use them as the measuring rod to compare the choice-worthiness of  $A$  in the world in which  $T_1$  is true and the choice-worthiness of  $A$  in the world in which  $T_2$  is true. Moreover, we have an answer to the question of grounds: the choice-worthiness difference between  $A$  and  $B$  on  $T_1$  is different from the difference in choice-worthiness between  $A$  and  $B$  on  $T_2$  in virtue of the fact that  $A$  and  $B$  instantiate different intrinsic choice-worthiness quantities in the world in which  $T_1$  is true than in the world in which  $T_2$  is true.

In general, on the account we suggest, if it is true that  $CW_i(A) - CW_i(B) = CW_j(C) - CW_j(D)$ , then it is true in virtue of the fact that the difference in the magnitude of the property of choice-worthiness that  $A$  instantiates and the magnitude of the property of choice-worthiness that  $B$  instantiates, in the epistemically possible world in which  $T_i$  is true, is the same as the difference in the magnitude of the property of choice-worthiness that  $C$  instantiates and the magnitude of the property of choice-worthiness that  $D$  instantiates, in the epistemically possible world in which  $T_j$  is true. In fact, as long as we know to take the following second-order claim at face-value, rather than analyze it in comparativist terms, we can state this claim in very natural language, namely: if  $CW_i(A) - CW_i(B) = CW_j(C) - CW_j(D)$  is true, then it is true in virtue of the fact that the difference between the choice-worthiness of  $A$  and the choice-worthiness of  $B$ , in the epistemically possible world in which  $T_i$  is true, is the same as the difference between the choice-worthiness of  $C$  and the choice-worthiness of  $D$ , in the epistemically possible world in which  $T_j$  is true.

Absolutism about choice-worthiness takes statements about choice-worthiness at face value: as ascribing an intrinsic property to an option. And once we allow the existence of necessarily existent choice-worthiness properties, then we have the resources to explain how intertheoretic comparisons are possible.<sup>35</sup> The absolutist about choice-worthiness mimics the

<sup>35</sup> One objection, raised to us by Christian Tarsney, is that there is a difficulty in moving from the fact that in the actual world, theories instantiate irreducibly monadic choice-worthiness properties to the conclusion that all moral theories in which the decision-maker has credence must be understood as committed to such properties. The former is a metaphysical claim; the latter a conceptual one. In other words: for our account to work, we need it to be the case that *all* the views in which the decision-maker has some credence must involve imputing irreducibly monadic choice-worthiness properties.

We believe that there are two ways in which one can respond to this worry. First, one could argue that the meaning of concepts like 'choice-worthiness' is determined by reference magnetism, and that theories that do not refer to such properties are simply changing the subject. However, though we find plausible the general idea that reference magnetism helps to determine the meaning of concepts, the claim that those with different metaethical views aren't even making coherent moral claims seems implausibly strong to us. So we prefer a second approach,

absolutist about mass in this respect: the absolutist about mass takes statements about the mass of objects at face value (as ascribing an intrinsic property of mass to an object), and then uses this to explain how inter-world mass relations are possible (as in the mass-doubled world case).

Before concluding this chapter, we note that it is not ad hoc to side with absolutism about choice-worthiness, rather than comparativism. This is for three reasons.

First, there are strong independent reasons that motivate absolutism. Not only can it explain intertheoretic comparisons of choice-worthiness, it can also explain other sorts of value-comparisons. For example, it can explain how we can make comparisons across *worlds*: we can understand 'x could have been better than x is' as saying that the value x has in the actual world is less than the value x has in some different possible world. Also, it can explain how we make comparisons of value across *time*: we understand 'x is better now than it used to be' as saying that the value x has now is greater than the value x used to have. Finally, it can explain comparisons of value between *mental attitudes* and the *world*: we understand 'x is better than I thought it was' as 'the value x has is greater than the value I thought x had'. These explanations are all intuitively simple. In contrast, it is unclear how the comparativist could offer equally intuitive and simple explanations of value-comparisons.

Second, we argued above that amplified theories are possible and that intertheoretic comparisons are clearly possible sometimes. Insofar as absolutism can give a natural and plausible explanation of that, whereas comparativism seemingly cannot, we have reason to prefer absolutism about choice-worthiness.

Third, the principal reason for rejecting absolutism about quantities of mass (and other scientific quantities) is a worry about needing to posit abstract entities such as properties in one's ontology. Whether or not this argument is successful in general, it is considerably weaker in the case at hand. In Chapter 7, we argue that moral uncertainty is inconsistent with non-cognitivism and, for the purpose of the project in this book, we must assume that error theory is false (otherwise there would be no subject matter for us

which is to distinguish metaethical uncertainty and normative uncertainty. Our account of how to make intertheoretic comparisons makes sense conditional on a particular metaethical view. We do not take ourselves to give an account of how to make intertheoretic comparisons across all metaethical views. In particular, insofar as the decision-maker should retain some credence in comparativism about choice-worthiness, any moral theories that are conditional on the comparativist view may be incomparable with moral theories that are conditional on absolutist views.

to investigate). That leaves us with some form of moral realism. And if we believe moral realism, then, though not an inconsistent combination of views, it certainly seems like an *odd* combination of views to be happy with the existence of moral facts, but to be sceptical of the existence of moral properties.<sup>36</sup>

### VIII. The Metaphysical and Epistemic Questions

Once we have accepted that the correct account of intertheoretic comparisons is non-structural, the problem of intertheoretic comparisons divides into two problems. The *metaphysical* problem is about what *grounds* intertheoretic comparisons.<sup>37</sup> That is: in virtue of what are intertheoretic comparisons true, when they are true? As we have argued, we think that the relations between abstract quantitative choice-worthiness properties ground intertheoretic comparisons.

But there is a further problem to be resolved. This is the *epistemic* problem: how can we tell which intertheoretic comparisons are true, and which are false?

Answering the first problem tells us about the nature of intertheoretic comparisons—what makes intertheoretic comparisons true. Answering the second problem would enable us, at least to some extent, to more confidently make intertheoretic comparisons: to more confidently know how two theories compare, when they do compare; and to more confidently know whether two theories are comparable at all.

In response to the epistemic problem, our view is rather deflationary. Because we endorse a universal scale account, we believe that, for any theory  $T_1$  and for any real number  $k$ , we can make sense of another theory  $T_2$  whose

<sup>36</sup> An additional objection, given to us by Christian Tarsney, is whether, on our view, it is possible to have an amplified credence distribution: that is, whether it is possible for there to be two decision-makers,  $D_1$  and  $D_2$ , that have all the same credences in all the same moral theories and intertheoretic comparison claims, except that all theories in which  $D_1$  has credence are amplified versions of the theories in which  $D_2$  has credence. If this is possible, it seems we then get into trouble—we now have to choose between an infinite number of credence distributions, and whichever one we pick we are basically guaranteed to be wrong.

We accept that amplified credence distributions are indeed a possibility on our account. Our response, here, is to appeal to a very weak form of reference magnetism: of all the possible amplifications of her credence distribution that she could have, the credence distribution she actually has is determined by what choice-worthiness properties are instantiated in the actual world. This is also what guarantees that at least some of her beliefs are (at least approximately) true.

<sup>37</sup> For discussion of the idea of grounding, see Kit Fine, 'Guide to Ground', in Fabrice Correia and Benjamin Schnieder (eds), *Metaphysical Grounding: Understanding the Structure of Reality*, Cambridge: Cambridge University Press, 2012, pp. 37–80.

choice-worthiness function is  $k$  times that of theory  $T_1$ . That is: every possible amplification of  $T_1$  is itself a distinct theory. So when we ask: 'How, if at all, do utilitarianism and this rights-based non-consequentialist theory compare?' we're really asking: 'Which, of the infinitely many different theories that have the same interval-scale structure as utilitarianism, and which, of the infinitely many different rights-based non-consequentialist theories, should we have most credence in?'

This means that the 'epistemic question' of which intertheoretic comparisons are true is really a question about how we ought to apportion our credences across different amplifications of a given class of equivalent interval-scale theories. And we believe that the methodology for answering that should be approximately the same as the methodology for first-order normative ethics in general: relying on intuitions about particular cases and appealing to more theoretical arguments.

To take an earlier example, consider Sophie, who initially believed partialist utilitarianism, but then became unsure between that view and the view according to which all persons have equal moral weight. The question about how to make intertheoretic comparisons between those two views reduced to the question of which, of all infinitely many theories within the class of classical utilitarian theories (including what we called kinship utilitarianism and Benthamite utilitarianism) she should come to have credence in. If she was moved to classical utilitarianism because it is a simpler theory, then it seems plausible that she should come to have most credence in Benthamite utilitarianism. If she was moved to classical utilitarianism by reflecting on the fact that there is a deep arbitrariness in whom she happens to have special relationships with, then it seems plausible that she should come to have most credence in kinship utilitarianism. Either way, we can explain why, as is intuitive, she should come to have most credence in one of those theories, rather than a different theory (according to which, perhaps, the value of distant strangers' welfare is 1 million times as great as it is on the partialist theory). Basic epistemic conservatism suggests that she should alter her beliefs as little as possible in order to accommodate new evidence (in this case, new arguments). Having partial belief in partialist utilitarianism, and partial belief in anything other than kinship or Benthamite utilitarianism, would be oddly incoherent.

If the account we have given is correct, this is an exciting development for first-order normative ethics. Moral theories, when they have been given, have really been *classes* of moral theories. And different views within this class can be more or less plausible than other views within this class. So

there may be scope to revisit old ethical theories, and assess which specific versions of those theories are most plausible.<sup>38</sup>

## Conclusion

In this chapter, we have argued against both sceptics and structuralists. Sceptics cannot account for the fact that we have intuitions about intertheoretic comparisons in many cases, and they have not provided a compelling general argument for their view. Structuralists cannot account for the ways in which aspects of theories other than their choice-worthiness functions seem to make a difference to how those theories should be normalized. We should therefore look for a theory of intertheoretic comparisons within the class of non-structural accounts. We defended an account analogous to Mundy's account of the metaphysics of natural quantities, arguing that intertheoretic comparisons are meaningful because of the relations between quantitative choice-worthiness properties.

Having completed our discussion of informational issues arising from MEC, let us next turn to two potential problems for our account.

<sup>38</sup> For example, Frances Kamm and Thomas Nagel claim that utilitarianism is implausible because it does not posit the existence of rights, and therefore that humans do not possess the value of dignity that can only be conferred by the possession of rights (Frances Kamm, 'Non-Consequentialism, the Person as an End-in-Itself, and the Significance of Status', *Philosophy & Public Affairs*, vol. 21, no. 4 (Autumn 1992), pp. 354–89; Thomas Nagel, 'The Value of Inviolability', in Paul Bloomfield (ed.), *Morality and Self-Interest*, New York: Oxford University Press, 2008, pp. 102–16). But neither Kamm nor Nagel distinguish between two different versions of utilitarianism. According to the first, no one has any rights, and so humans are indeed of less value. According to the second, people do have rights not to be killed (for example), but they also have equally strong rights to be saved. Both have the same interval-scale structure. But, according to the latter form of utilitarianism, humans *do* have the value of dignity that can only be conferred by having rights. So Kamm's and Nagel's argument would not go through. See Shelly Kagan, *The Limits of Morality*, Oxford: Oxford University Press, 1989, ch. 3.