11 Moral uncertainty and human embryo experimentation

Graham Oddie

Moral dilemmas can arise from uncertainty, including uncertainty of the real values involved. One interesting example of this is that of experimentation on human embryos and foetuses. If human embryos or foetuses (henceforth *embryos*) have a moral status similar to that of human persons then there will be severe constraints on what may be done to them. If embryos have a moral status similar to that of other small clusters of cells then constraints will be motivated largely by consideration for the persons into whom the embryos may develop. If the truth lies somewhere between these two extremes, the embryo having neither the full moral weight of persons, nor a completely negligible moral weight, then different kinds of constraints will be appropriate.[1] So on the face of it, in order to know what kinds of experiments, if any, we are morally justified in performing on embryos we have to know what the moral weight of the embryo is. But then an impasse threatens. For it seems implausible to many to suppose that we have settled the issue of the moral status of human embryos. It is the purpose of this chapter to show that such moral uncertainty need not make rational moral justification impossible. Section 1 develops a framework, applicable to cases of moral uncertainty, which distinguishes between what is morally right/wrong, and what is morally justified/unjustified. In section 2 this framework is applied to the case of embryo experimentation. The upshot is that the kinds of goods which would justify performing experiments on human embryos are of the same order of value as goods which would justify comparable experiments on non-consenting persons.

1 A framework for moral uncertainty

1.1 *Three cases of uncertainty*

Case 1 Your infant son is seriously, but not terminally, ill. You have two drugs: drug **A** would effect a partial although not complete cure and this you know for sure; drug **B** is rather unusual in that its efficacy

depends on chance factors at the molecular level. With your particular child it has a 2/3 chance of the right reactions taking place, thereby curing him completely, and a 1/3 chance of the wrong reactions, killing him. These are genuine physical chances, not a matter of your ignorance. But of course, even though you know all this, you cannot know in advance what the outcome will be if drug **B** is administered. Genuine chanciness implies uncertainty, while uncertainty does not necessarily imply genuine chanciness.

Case 2 Again your child is seriously, but not terminally, ill. This time you have three drugs: drug **A** as above, and drugs **C** and **D**. One of these drugs would (with chance 1) completely cure him if administered, the other would kill him (again with chance 1) if administered, but you do not know which is the killer and which the cure. To give him drug **C** would, in fact, be an act of curing him. To give him drug **D** would be to kill him. So you are uncertain which of the two acts would be the act of killing, and which the act of curing. However, you have evidence that gives you a 2/3 degree of confidence that it is drug **C** which completely cures, and a 1/3 degree of confidence that it is drug **D**.

Case 3 Your pet monkey is seriously ill. Drug **E** will cure the animal but will leave it with some discomfort and pain for the rest of its life. Drug **F** will end the animal's life painlessly. You know this, but you are unsure about the relative merits of painlessly annihilating a monkey and keeping it alive impaired. However, you have some views on the matter. You think it more likely that the monkey has a moral status akin to that of a human infant, rather than to that of a cockroach.

We are often in a position of ignorance or uncertainty about the nature and value of our actions, or of their outcomes, as cases such as these illustrate. This is obviously a problem for consequentialist moral theories. If the morally right action to perform is that with the best consequences, then to know what's right will typically require extensive knowledge not only of the value of the various possible outcomes of one's actions, but also of the causal structure of the world. In case 2 above the best action with the best consequences *in fact* is **C**, since that is the one which effects the total cure, although the agent does not know this, because he does not know the causal structure of the situation. But ignorance of causal structure is not the only generator of uncertainty. In case 1, there is uncertainty, but this time stemming not from ignorance of causal factors, but from objectively chancy factors at the micro-level. In case 3, on the other hand, you are clear about the objective chances, but uncertain about

both the relative value of the monkey's life, impaired to that degree, and of painlessly ending its life.

While the problem of uncertainty is obvious for the consequentialist, deontological theories are not immune. Most kinds of actions of moral interest involve necessary conditions which go beyond the agent's behaviour, conditions which must take place if the behaviour is to constitute an act of a specified kind. Take killing and curing, for example. One cannot kill a person without him dying. One cannot cure a person without her returning to health. The dying and the returning are not part of the agent's direct behaviour. Rather the agent does something (like administering drug **C** or administering drug **D**) over which he does have more or less direct control, which has as its causal effects the dying, or the returning to health. So even if an agent can be sure that it is wrong to kill a person and right to cure him, he may find himself in a situation in which he cannot be sure which, if any, of the actions over which he has direct control will constitute an act of killing or an act of curing.

The deontologist might be tempted to respond that it is not *killing*, but rather the *intention* to kill which is morally wrong. According to this line, killings would be wrong because each and every one involves the intention to kill, and murderous intentions that do not result in deaths would be morally equivalent to those that do. Since the intention is (presumably) something of which the agent can be directly aware, and which he can (presumably) directly control, the problem of uncertainty would not arise.[2]

By declaring all the actions of moral significance to be such mental acts the deontologist may escape the problem of uncertainty – but only at enormous cost. For it seems clear that at least part of what is wrong with killing someone is that someone is killed. The wrongness of murderous intentions is at least partly to be explained by the fact that in general such intentions tend to have undesirable effects beyond the agent's own mind.

1.2 *A decision-theoretic framework*

Suppose, firstly, that the moral *rightness* or *wrongness* of an act is determined by the overall goodness or value of the act, compared with the overall goodness or value of its alternatives. For example, an act may be right only if there is no alternative option which is more valuable. Suppose, secondly, that the value of an act depends in some way on features of which you may have only limited knowledge. For example, it may be that the value of an act depends partly on the value of the actual or probable consequences of the act. Given these two assumptions you may not always know what option is the morally right one because you may

not always know the precise comparative values of your options. The best you can do is to *estimate* the values of your various options on the basis of the information and beliefs you do have, and it is plausible that what you are morally *justified* in selecting is the act with the highest estimated value, by your lights.

So our two assumptions yield a natural distinction between rightness/wrongness (governed by real value), and the moral justifiability/unjustifiability (governed by estimated value in the light of your information). For example, in case 2 it would be natural to *estimate* the value of drug **A** as greater than that of either drugs **C** or **D**, since by your lights both of the latter bestow a significant probability on something of very great disvalue: killing your child.

By subordinating questions of rightness to those of value this framework may have a consequentialist ring to it. But because of its compatibility with diverse accounts of value it is applicable to an extremely broad range of moral theories, including some which it might be natural to classify as deontological. For example, it is perfectly compatible with the assigning of value or disvalue to an agent's intentions, or the particular character of the agent's performances, the kind of people we become through the performance of a certain kind of act, to the exemplification of the virtues, and so on.

But exactly how does the value of an act depend on features like its consequences? According to traditional consequentialism the value of an act is determined by how much value it would realise. But a serious problem arises because the decisive quantity, *the amount of goodness act A will/would realise*, may be undefined. Consider case 1. With the option of prescribing the chancy drug **B** there are two different possible outcomes, with very different values. If in fact the drug is administered then one could say that the value of doing so is the value of the actual outcome, either death or cure, whichever it happens to be. But suppose drug **B** is *not* administered. In that case there is simply *no answer* to the question 'How much value *would* have been realised had drug **B** been administered?' because it is not true of either outcome that it *would* have ensued if the drug had been administered. It might have produced a cure (chance 2/3) and it might have produced death (with chance 1/3).

It is here that we can take a leaf from the decision-theorist's book: the quantity which determines the value of the action is the amount of value one could (in a technical sense) *expect* the option to realise if one knew the objective chances. How much value is an action likely to realise? It is natural to measure this by its *expected* value.[3] The expected value of an action A is the weighted average of the real values of the possible outcomes of A, where each outcome is weighted by its *probability* given A.

There is thus an analytical link between probability and value: an action A has greater (expected) value the more valuable are its possible outcomes, or the more probable are its valuable outcomes.

Expected value admits of various interpretations, depending on whether the probability and value functions are either subjective or objective. Where both the probabilities and the values are objective, then we have *objective expected value* and it is this quantity which most obviously suggests itself as the determiner of an action's objective *rightness*. On this view, the morally right action is the one which would, objectively, most conduce to the good, regardless of the agent's beliefs about the good or about the causal structure of the world.

The objective chance of an outcome given A is not an epistemic matter: it is the real propensity A would bestow on that possible outcome if A were to be performed. Nor is objective value an epistemic matter. In case 1, prescribing drug B has two possible outcomes, *complete cure and death*, with respective objective chances of 2/3 and 1/3 (regardless of any particular agent's subjective probabilities) and a large value difference (regardless of any agent's desires or beliefs about value).

Let this value difference between the two outcomes (value of curing–value of killing) be c. It is only differences in value which are significant in the comparison of actions, so we can conventionally set the origin of our scale, 0, at the value of killing, and the value of curing at c. The objective expected value of administering drug **B** is the chance-weighted average of the value of the killing (0) and the value of a complete cure (c):

objective expected value of prescribing drug **B**
= chance of killing given that drug **B** is prescribed × value of killing + chance of a complete cure given that drug **B** is prescribed × value of a complete cure
$= \frac{1}{3}0 + \frac{2}{3}c = \frac{2}{3}c.$

In the general case where the respective chances are p and $(1-p)$ for curing and killing, the objective expected value of administering **B** would be pc.

On the other hand a partial cure has some positive value, a value which is clearly not as great as a complete cure, but is vastly better than death: say rc, where r is a real number between 0 and 1. r measures how close (on a scale of 0 to 1) the value of a partial cure is to the value of a complete cure, and how far it is from the (dis)value of killing. Since a partial cure is much closer to a total cure than to killing, r will be rather close to 1. Prescribing drug **A** ensures (with chance 1) a partial cure. The objective expected value of prescribing drug **A** is thus just rc, and this exceeds the expected value of prescribing drug **B**, provided that r is greater than p

(= the probability that **B** will produce a cure). In case 1, if the agent ought to choose by objective value, then drug **A** is the right drug to prescribe.

Note that one might motivate the decision to prescribe drug **A** on grounds of *safety first*: that because drug **B** yields *some* risk of death, and drug **A** does not yield any, drug **A** is to be preferred. Lying behind this rationale is the principle of avoiding the worst possible outcome. (Minimise maximal possible loss.) However, the principle to which I am appealing (maximise expected value) is quite different. It tells us that a risk of death may be acceptable, *provided the chance of death is small enough to be swamped by the probability of a good outcome*: provided, that is, r is greater than p. But given that the value difference between a complete cure and death is enormously greater than that between a complete cure and a partial cure, the chance of death would have to be compensatingly very small for the risk to be worthwhile.[4]

The kind of uncertainty faced in case 1, where one drug bestows different objective chances on two possible outcomes, is analogous to the kind of uncertainty faced in case 2. But there is a difference: in case 2 what you don't know are the objective chances that your options bestow on the possible outcomes. You cannot directly apply the injunction to maximise objective expected value, since you don't know what the objective chances are. The injunction does not give you an effective selection procedure, one which is accessible to your mind. If your goal is to maximise objective value, but you don't know where maximal objective value lies, what effective selection procedure are you justified in using?

In place of employing the actual but epistemically inaccessible chances in the calculation of value, you are justified in employing your *best estimate* of these. While you may not know the objective chances, you have some information about them or at least some beliefs, and this information gives you a guide to the objective chances: possibly not a very good guide, but the best one you have. Call these estimates of the objective chances your *subjective* probabilities. If you know the actual value of outcomes, but are uncertain of the chances which your actions bestow on those outcomes, your *best estimate* of objective expected value combines your subjective probabilities with the real objective values. Subjective expected value is thus the best estimate you have of the objective expected value of your options.[5]

In case 1 substituting your subjective probabilities for the objective chances yields:

> estimated value of drug =
> subjective probability of partial cure given drug × real value of partial cure

+ subjective probability of complete cure given drug × real value of complete cure
+ subjective probability of death given drug × real (dis)value of death

It is not hard to calculate the results of estimated value (and of objective value):

	estimated value	objective value
drug **A**	rc	rc
drug **C**	$\frac{1}{3}0+\frac{2}{3}c=\frac{2}{3}c$	c
drug **D**	$\frac{2}{3}0+\frac{1}{3}c=\frac{1}{3}c$	0

The estimated values are not uniformly identical to the objective expected values. But they are still your best guide to the objective values of the acts given your epistemic situation. I submit that an agent is *epistemically justified* in ranking the options according to their estimated values, and *morally justified* in choosing according to that ranking: 'morally', because of the use of real values; 'justified', because of the use of the epistemic probability measure. In case 2 drug **C** is the (objectively) right one to prescribe (as given by objective value) since it effects the complete cure – but it is not the morally justified option. The morally justified option, the one with highest estimated value in the light of your information, is drug **A**. And that is, what we would intuitively expect. Prescribing drug **C** in case 2 would be morally blameworthy, because morally unjustifiable, even though it is the objectively right thing to do.

The interesting point here is you may not only be uncertain of the objective value by virtue of uncertainty of the objective chances of outcomes. Your uncertainty may extend to the *values*. In case 3 you know what effects the drugs will have on your pet monkey, but you may well be uncertain about the values and disvalues involved. Is it really better for such an animal to live an impaired life than for that life to be painlessly ended? Perhaps there is no answer to this question, because there is no such thing as objective value. But even if there is a correct answer you don't know what it is.

I will argue that the moral uncertainty of embryo experimentation is rather like case 3. As yet we may feel we have no entirely satisfactory solution to the problem of the moral status of the embryo, and, further, this particular lacuna is not culpable ignorance. It would be, say, if by a little bit more effort on our part we could clear the matter up to the satisfaction of rational creatures apprised of all the facts accessible to us. But that seems implausible. And even if it were true, there could be other cases of non-culpable ignorance, or uncertainty, of the objective values of

outcomes. Indeed, if values are objective *in some sense*, if there can be a truth of the matter as to what value we ought to attribute to outcomes, then we are almost obliged to conclude that non-culpable uncertainty of real values is a rather common phenomenon.

Given a case in which uncertainty extends to the values themselves, what choices is an agent morally justified in making? It is not enough to apply the proposal above, which employs the agent's estimates of chance together with real values of outcomes, as it stands. As in case 2, where the real chances are unknown and what is morally justified is determined by your best estimates of the chances, so in case 3 what is morally justified should be determined not by the uncertain objective values, but by your best estimates of these. If you are *rationally justified* in estimating the real value of B to be greater than that of A then you are morally justified in performing B rather than A, even if B is in fact objectively worse, and hence the morally wrong act.

Sense can be made of this simple idea within our decision-theoretic framework provided uncertainty of the objective values of outcomes is treated as formally tantamount to facing a number of different possible outcomes with differing objective values. Uncertainty of the real value of a particular outcome generates a set of possible (real) values of that outcome. Call these *value possibilities*. You may well have views about which value possibilities are more likely (in the purely subjective sense) than others. Idealising we can imagine that your views about value can, like your views about the natural causal facts, be represented by degrees of subjective probability.

For example, in the case of your pet monkey you may think there are two value possibilities: the first, that killing the monkey is morally equivalent to killing, say, a human infant; the second, that it is morally equivalent to killing, say, a cockroach. Suppose that you have reasons (whatever their source) for inclining to the former view (2/3 subjective probability) rather than the latter (1/3 subjective probability). It is now fairly obvious that you can estimate the value of the various options taking into account the range of value possibilities and the relative likelihood of those possibilities. In general, then, you are *morally justified in choosing to do A* just in case A has maximal estimated value.

There is clearly a proviso to be added here: that your ignorance or uncertainty is not culpable. There is more to be said about this, for the addition of the proviso threatens circularity. In order to specify what an agent is justified in doing, it has to be specified what he is justified in believing. A full elaboration of this account would indeed have to show that this apparent circularity is not vicious, but for our purposes we need only work with an intuitive grasp of culpable ignorance. Intuitively, an

agent is culpably ignorant if he could have secured more accurate beliefs relevant to his decision problem, and he was aware of his ability to do so.

This decision-theoretic account thus has a solution to the problem of morally justifying actions in the absence of complete certainty about the objective values of outcomes, and it is to an interesting application of the framework that we now turn.

2 An application: the case of embryo experimentation

2.1 *Lethal experimentation on non-consenting persons*

It is widely agreed that embryos used for experimental purposes ought not to be allowed to develop into fully fledged persons. In general it is not known, in advance, quite how an experiment would affect the embryo's development. Even quite minor interventions could have dramatically detrimental effects on the person into whom the embryo would develop if allowed. Thus *The Warnock Report*, for example, recommends that no embryo used for experimental purposes be allowed to develop beyond fourteen days.[6] If this recommendation were accepted then all experimentation on embryos would be lethal, and the question of the permissibility of embryo experimentation would reduce to that of the permissibility of lethal embryo experimentation. For simplicity I will deal with the case of lethal experimentation first, and consider later whether embryo experimentation could be justified if the Warnock recommendation were rejected.

It is important to clarify the concept of a lethal experiment. An *experiment* consists in forcing an object, or system of objects, into a state it would not, or might not, otherwise enter, in order to find out what will happen to it. Experimentation thus differs from observation in that the former involves intervention and the latter does not. Experimentation also differs from therapy in that the primary end of experimentation is knowledge, whereas the primary end of therapy is healing. A *lethal experiment* is one performed on a living organism with *either* the intention that the experimental state lead directly or indirectly to the death of the organism, and it does so lead, *or* a reasonable belief that such an outcome is the likely result of the experiment, and it is.

Lethal experiments on non-consenting persons are almost universally thought to be seriously wrong. Moreover, the wrongness of such experiments is not thought to consist solely in their lethality. Many killings which are wrong do not seem as wrong as lethal experiments. Some, at least, of the carnage of World War Two was wrong, and judged to be so. But the lethal experiments routinely performed by Joseph Mengele on

non-consenting persons were widely judged to be of a different order of wrongness altogether. There will be different explanations of this judgement on different moral theories, but it is not necessary here to adjudicate between them. This judgement on lethal experiments can be taken here as a *prima facie* moral datum which any moral theory must attempt to accommodate. The judgement could be overturned (corrected) by an otherwise extremely powerful and attractive moral theory but, other things being equal, a moral theory which delivers this judgement is to be preferred to one which does not.

A programme of lethal experiments on non-consenting persons might well be a promising source of very valuable information. Such information might have intrinsic scientific value (if indeed information has intrinsic value), or it might yield medical procedures which would benefit many other people enormously. The fact is, however, that such programmes are not usually put forward for research funds. Or, if they are, they are turned down. Or, if they are not turned down, the public does not get to hear about them. Or, if the public does get to hear about them, there is general moral outrage.[7] An excellent decision-theoretic explanation for this collection of facts is that the disvalue of a programme of lethal experiments on non-consenting persons is judged to outweigh heavily the value of the kinds of goods which such a programme might be expected to yield.[8] There may be possible circumstances which would justify a programme of such experiments, and it may be that in such circumstances people would agree that they were justified, even if no one were willing to volunteer. For example, they might be justified if it appeared likely that the only alternative was the extinction of the human race, or its survival under totally unacceptable conditions. But all that is required for our purposes here is that such circumstances do not often turn up.[9]

2.2 *The moral status of the embryo*

If the abortion debate has made one thing clear it is the extraordinary difficulty of achieving some kind of rational consensus on the moral status of human embryos. Note that while the issues of abortion and experimentation both centre on the problem of the embryo's moral status it should not be assumed that positions on abortion will simply coincide with positions on embryo experimentation. Those who think that abortion is morally permissible are not compelled to accept the permissibility of lethal experimentation. Firstly, embryos which are used for experimental purposes will normally be sustained in an artificial environment, outside the womb, and so the question of a woman's rights with respect to her own

body need not enter the picture.[10] Secondly, if section 1 is on the right lines, lethal experiments can reasonably be regarded as worse than ordinary killings. It can be permissible (or obligatory) to kill a being without it being permissible (or obligatory) to lethally experiment on it. Thus we can afford to sideline the problem of the moral justifiability of abortion here, though we cannot afford to ignore whatever that debate has taught us about the moral status of the embryo.

Wertheimer argued, in 'Understanding the abortion argument' that the prospects for a satisfactory, rational resolution of the issue of the embryo's moral status were poor, and the subsequent development of the debate is evidence in favour of Wertheimer's claim.[11] There is no doubt that the issues are much clearer now than they were twenty years ago. For example, we now know not to confuse the concept of a human being with that of a person, or that of a human person.[12] It is fairly well established that purely biological characteristics are not morally relevant.[13] Much excellent work has been done to clarify the moral relevance of potential personhood. But despite these significant gains in understanding, the debate seems not to have reached a satisfactory conclusion, let alone any kind of consensus.

There will, of course, be those on both sides of the debate who will disagree strongly with this claim. Instead of attempting to convince them by rehearsing the old arguments, or attempting to construct new arguments for either side, I will take Wertheimer's claim as a premise, and see what follows from it. Those who are reluctant to embrace fully the consequences of either of the rival positions must acknowledge that we are uncertain about the moral status of the embryo. Moreover, if Wertheimer is right, this ignorance does not stem from a failure to search thoroughly for the truth, but from the nature of the problem. What are we justified in doing with embryos given an acknowledgement of that uncertainty?

The application of the decision-theoretic framework to the case of lethal embryo experimentation will be developed in three stages, at each stage relaxing an important simplifying assumption.

2.3 *The first version: two possibilities for moral weight*

The basic structure of the argument is simple. We face a choice between experimenting on embryos and not experimenting. Simplifying, either the embryo does have the same moral status as a person or it has no moral weight at all. (That the embryo has the same moral weight as a person is not, of course, equivalent to saying that it *is* a person.) This generates four possibilities:

	Embryo has moral weight of person	Embryo has no moral weight
Experiment	First possibility	Second possibility
Don't experiment	Third possibility	Fourth possibility

On the assumption most favourable to experimentation, experimenting will deliver valuable goods not otherwise obtainable. Let us agglomerate these and dub them, *the good*: of value g. So this good is obtained in both the first and second possibilities, but not in the third and fourth. On the other hand, by experimenting there is the possibility of something decidedly lacking in value: lethally experimenting on a being with the same moral weight as a person. Call this, *the evil*: of value l. It is crucial to note that the evil is *equivalent to the evil of lethally experimenting on a non-consenting person.* This evil arises only given *both* that we experiment *and* that embryos have the same moral weight as persons: the first possibility. So the good arises only in possibilities one and two, and the evil only in possibility one. Let us assume that the good and the evil are commensurable and tradable. In other words, the value of the first possibility is the sum of the good and the evil, $g-l$, whereas the value of the second possibility is just the good alone: g. It is only value differences that count in ranking possibilities and, since the third and fourth possibilities involve neither the good nor the evil, we can set the value of both at 0. Thus our table of value looks like this:

	Embryo has moral weight of person	Embryo has no moral weight
Experiment	$g-l$	g
Don't experiment	0	0

The salient feature of this table is that, provided the argument of section 2.1 is right, it is strongly asymmetrical. In normal cases l will be very large in comparison with g. It is this asymmetry which does most of the work in what ensues.

To rank the two possible courses of actions we compare their subjective expected values. Let p be one's subjective probability that the embryo has moral weight, and $(1-p)$ the probability that it does not. Then the expected value of experimenting is:

probability of the first possibility given experimentation × value of the first possibility
+ probability of the second possibility given experimentation × value of the second possibility
$= p(g-l)+(1-p)g$

$= g - pl$
= (value of the good) minus (the disvalue of the evil × probability of the embryo having moral weight).

The expected value of not experimenting is just 0 ($= p0 + (1-p)0$). So for experimentation to be preferable, the expected value of experimenting must be greater than that of not experimenting. And that will hold just in case g is greater than pl. That is, the value of the good must be greater than the disvalue of the evil × probability of the embryo having moral weight.

We are making a number of simplifying assumptions which favour the case for experimentation. Firstly, the costs in terms of resources which such experiments must involve are ignored. Secondly, even if embryos do have the same moral status as persons we assume that the goods obtained by experimenting retain their value despite their etiology. (This may seem innocuous, but it could easily be challenged. Consider the moral indignation when it was revealed a number of years ago that scientists were using data Mengele had derived from his lethal experiments on non-consenting persons.) Thirdly, it is certain that experimenting will yield the good and that failing to experiment will deprive us of the good.

I have argued (section 2.1) that the magnitude of the evil at issue (lethally experimenting on non-consenting persons) is much greater than that of the good we could hope to derive thereby. So for the inequality to hold the probability of the embryo having moral weight must be vanishingly small. But given uncertainty as to the embryo's moral weight (section 2.2) that is false.

We can approach the result in another way without relying on the conclusion of section 2.1. For lethal experimentation on non-consenting persons (disvalue $= l$) to be justified by the certain prospect of, say, a cure for leukemia (value $= g$), we would have to have $g > l$. But given our ignorance of the moral status of the embryo (say, $p = 1/2$) to justify lethal experimentation on embryos we would have to have: $2g > l$. But then g (the value of the cure) and l (the disvalue of lethal experimentation) are not of a different order of magnitude: g is within reach of l. Suppose that the value of a cure for leukemia is roughly that of a cure for AIDS, and the value of having both is roughly twice that of having one such cure. Then if we judge it justifiable to experiment lethally on embryos to obtain one of these cures, the certitude of both cures would justify a programme of lethal experimentation on non-consenting persons. Looked at in this way the argument shows that a positive attitude to embryo experimentation calls for a rather drastic revision in our judgement of the relative merits of lethally experimenting on non-consenting persons.[14]

2.4 *The second version: continuity of moral weight and pure agnosticism*

Beings do not fall into just two categories, those with the full moral weight of persons, and those with zero moral weight. Rather, moral weight is (at least in principle) a continuous magnitude.[15] The first version of the argument is thus simplified, but the simplification is not vicious.

We may assume that the moral weight of the embryo (w), compared with that of a human person, lies between 0 and 1 (inclusive), and that the disvalue of lethally experimenting on embryos is wl. The moral weight of the embryo is thus fixed by the intrinsic disvalue of lethally experimenting on embryos compared with the intrinsic disvalue of lethally experimenting on non-consenting persons. It is the ratio of the former to the latter. We thus rule out the possibility that embryos are worth more than ordinary human persons ($w > 1$) and the possibility that lethally experimenting on them could be *intrinsically* valuable ($w < 0$). We are now dealing with a continuous magnitude, and the discrete probability assignment to the two value possibilities gives way to a probability density over the possible values of the magnitude w. Given your particular probability density we can calculate the *expectation* of w (E_w) which is, roughly speaking, what you could expect w to be, given your particular beliefs. We could call E_w the embryo's *estimated* moral weight, and it turns out that estimated moral weight plays exactly the same role in this second version as does p (the probability that the embryo has the full moral weight of a person) in the first version. That is to say, it is easy to show that experimentation is justifiably preferred to the alternative just in case $E_w < g/l$.

Pure agnosticism over the precise location of the value of w in the unit interval is tantamount to a principle of indifference: that the probability distribution over the possible weights is perfectly even. This yields the result that E_w, the estimated moral weight of the embryo, is 1/2. Thus the assumption of pure agnosticism about the embryo's moral weight, yields precisely the same result as an equal probability assignment to the two extremes of the spectrum. As before experimentation is morally justifiable just in case: $2g > l$.

2.5 *The third version: relaxing agnosticism*

Pure agnostics thus have grounds for resisting lethal embryo experimentation comparable to the grounds (if any) they have for resisting lethal experimentation on non-consenting persons. While many would agree with the pure agnostic that we do not known for certain what the moral weight of the embryo is, they would claim that the balance of probabilities

is definitely in favour of a low weight rather than a high weight. They might claim that the arguments against a high weighting are not absolutely compelling but that they are nevertheless more powerful than the opposing arguments. So we may be relatively confident that the moral weight of the embryo is small rather than large.

The assumption of relative confidence could be spelt out in a number of ways. It would have to entail that the estimated moral weight of the embryo is less than 1/2. But even quite a small estimated weight, say 1/10, would still generate a moral problem for experimentation. It would entail that if lethal experimentation on embryos is morally justified to obtain a good g, then lethal experimentation on non-consenting persons is justified to obtain goods ten times the value of g.[16]

2.6 *Non-lethal experimentation*

An advocate of embryo experimentation might well take this argument to be a *reductio* of *The Warnock Report*'s recommendation that embryo experimentation be lethal. The response might be that embryos used for experimental purposes ought to be given a reasonable chance to survive and develop. But suppose that a certain kind of experiment on an embryo is advocated. The point of the experiment is to find out what will happen to the embryo when it is forced into the experimental state. If we knew in advance what would happen to the embryo there would be no point in subjecting it to the experiment. Thus if the experiment is to be of value we would have to be reasonably ignorant, in advance, of the likely effects on the embryo of the experimental state. Thus any such experiment would involve a risk of harm, possibly a considerable one, to the person into whom the embryo would develop. Even if we put to one side the issue of the embryo's moral status, non-lethal experimentation would be subject to the same constraints appropriate to experimentation on non-consenting persons. Thus non-lethal experimentation on human embryos is likewise justified only if the goods to be obtained are comparable to those which would justify similarly risky experiments on non-consenting persons.

3 **Implications**

The argument establishes that, given value uncertainty, it is morally justified to perform lethal (or risky) experiments on embryos only if it is morally justified to perform lethal (or risky) experiments on non-consenting persons to obtain comparable goods. Apart from that of value uncertainty, the assumptions of the final version of the argument are all either relatively uncontroversial or they favour embryo experimentation.

Of course, any argument can be regarded as a challenge to its premises. Thus those committed to the justifiability of experimentation may well take this to establish what they have assumed all along: that embryos are of negligible moral weight. Or they may even begin to suspect that the blanket proscription against lethal experimentation on non-consenting persons is itself unjustified. But what the argument certainly shows is that the strongly asymmetrical value structure of this particular moral dilemma imposes the onus of proof very heavily on one side of the debate. Moderate agnostics on the question are forced to take sides with those who claim that embryos have the same moral status as human persons. Thus even if experimentation on human embryos happens to be morally *right*, it will be difficult to justify morally. And, of course, an analogous conclusion will apply to any moral problem in which the possible value outcomes exhibit the same strong asymmetry.

ACKNOWLEDGEMENTS

Earlier versions of this chapter were read at seminars at Otago University and Hebrew University and at a meeting of the British Society for the Philosophy of Science in London. I would like to thank all of those who have made comments and criticisms: particularly Grant Gillett, Frank Jackson, Hugh Mellor, Roy Perrett and Philip Pettit.

NOTES

1 *The Warnock Report* (1984); that is, the *Report of the Committee of Inquiry into Human Fertilisation and Embryology*. (London: Her Majesty's Stationery Office), endorses an intermediate position of this kind. On the basis of this intermediate status the report argues (pp. 64–5) that 'embryo research should not be totally prohibited. We do not want to see a situation in which human embryos are frivolously or unnecessarily used in research ...', but '... continued research is essential, if advances in treatment and medical knowledge are to combine.'

2 This is one interpretation of Kant's theory of the good will: see 'The good will and its results' in *Groundwork of the Metaphysics of Morals* edn H. J. Paton (1972) *The Moral Law* (London: Hutchinson) p. 60.

3 See G. Oddie and P. Milne (forthcoming) 'Act and value: expectations and the representability of moral theories' in *Theoria* for the thesis that acts should be evaluated by the expected value of the various possible outcomes with which they are compatible.

4 A thorough application of the safety first principle would make our lives intolerable. See J. C. Harsanyi (1976) 'Can the maximin principle serve as a basis for morality? A critique of John Rawl's theory', in *Essays on Ethics, Social Behaviour, and Scientific Explanation* (Dordrecht: Reidel) p. 37ff.

5 See P. Menzies and G. Oddie (1992) 'An objectivist's guide to subjective value'

in *Ethics* 102: 3 pp. 512–33 for a full defense of this position. This draws on earlier work by Frank Jackson (1986) in 'A probabilistic approach to moral responsibility' *Proceedings of the Seventh International Congress of Logic, Methodology, and Philosophy of Science*, R. Barcan Marcus *et al.* (eds.) (Amsterdam: North-Holland Publishing Co.).

6 *The Warnock Report* (1984) p. 66.

7 J. Katz's (1972) *Experimentation with Human Beings* (New York: Russell Sage Foundation) is a fascinating compendium of material on the subject. Particularly interesting in this connection is the Jewish Chronic Diseases Hospital Case (1963) in which chronically (though not terminally) ill patients were, without their informed consent, injected with live cancer cells to test their immunological response. (Opinion is divided over whether this could have induced a serious malignancy.) These experiments were not lethal according to our definition (none of the patients died as a result of them, and probably none was intended to die) but they certainly bear a resemblance to lethal experiments, and the episode, as recorded in the documents assembled by Katz, does bear out our four-fold 'thesis' about such experiments. (1) The programme was not submitted for proper scrutiny by the appropriate authorities before it was implemented. (2) If it had been it seems it would have been rejected. (3) Had it not been for a patient's complaint the episode might never have been exposed. (4) And once it was exposed public reaction against it was strong.

8 One way of accommodating such a judgement is sketched by A. Sen (1982) in his 'Rights and agency', *Philosophy and Public Affairs*, 11, no. 1, pp. 3–39. A 'goal rights system' is one in which 'fulfilment and non-realisation of rights are included among the goals, incorporated in the evaluation of states of affairs, and then applied to the choice of actions through consequential links'. (p. 13) A much more systematic framework which shows that this is always theoretically possible is presented in Oddie and Milne (forthcoming).

9 A non-consequentialist might give a different explanation. For example, it might be argued that persons have a *right* not to be lethally experimented upon no matter what goods might be derived from such experiments. This can be assimilated to the decision-theoretic explanation by assigning a disvalue to lethal experiments which is so large that no feasible amassing of the sorts of goods which could be derived from such experiments could ever compensate. See Oddie and Milne (forthcoming) for the details.

10 J. J. Thomson (1971) has argued in 'A defense of abortion', *Philosophy and Public Affairs*, 1, no. 1, pp. 47–66 that even if the foetus (or the embryo) has the full moral weight of a person it is not always wrong to abort it. However, Thomson acknowledges that her argument shows only that the woman has the right to dislodge the foetus or embryo, not the right to ensure its death, let alone the right to perform lethal experiments on it. A Thomson-type argument could not establish the legitimacy of lethal experimentation on doomed foetuses or embryos, unless we add, as a general principle, that a doomed organism is fair game for lethal experimenters. And that, in general, is clearly unacceptable.

11 R. Wertheimer, (1971) 'Understanding the abortion argument', *Philosophy and Public Affairs* 1, no. 1, pp. 67–95.

12 For a clear account of the distinction and a survey of the arguments see M. Tooley (1983) *Abortion and Infanticide* (Oxford University Press), pp. 50–86.
13 *Ibid.*, pp. 61–77.
14 A related argument has been put forward by Lachlan Chipman in a short article 'By whom begot: IVF, law making and moral uncertainty', *Quadrant*, 28, no. 9, pp. 16–17. However, Chipman inclines to the *safety first* principle of minimising maximal loss. Consider, 'Killing people whose lives might be saved is so abhorrent, that the fact that there is a chance, even a remote chance, that our conduct would produce that result is enough to make it unacceptable'. (p. 17) What is all important is the *degree* of remoteness of the chance. Perhaps there is a remote (*very* remote) chance that potatoes have the same moral weight as persons. This should not stop us eating chips, for the simple reason that the chance is so *very* remote, and this is precisely what the principle of expected value delivers. (This also answers an attack by Wertheimer (1971) 'Understanding the abortion argument', *Philosophy and Public Affairs* 1, no. 1, p. 76, on a simple minded version of the argument from uncertainty.)
15 An interesting discussion of the idea that the physiological continuity from blastocyst to adult human yields some kind of continuity of moral weight is to be found in N. Poplawski and G. Gillett (1991) 'Ethics and embryos', *The Journal of Medical Ethics* June. Poplawski and Gillett argue that the value which attaches to a human life as a temporal whole must also attach (at least partly) to its temporal stages, and hence even to the early stages of such a life. This is compatible with (although does not entail) the possibility that the embryo has a significant value, albeit not the same value as a fully functioning person.
16 Also, if the estimated moral weight of the embryo/foetus increases with its age we get the intuitively correct result that the difficulty of justification also increases with age.