RAHUL KUMAR

PERMISSIBLE KILLING AND THE IRRELEVANCE OF BEING HUMAN

(Received 11 July 2006; accepted in revised form 11 July 2006)

ABSTRACT. This is a review essay of Jeff McMahan's recent book *The Ethics of Killing: Problems at the Margins of Life* (OUP: 2002). In the first part, I lay out the central features of McMahan's account of the wrongness of killing and its implications for when it is permissible to kill. In the second part of the essay, I argue that we ought not to accept McMahan's rejection of species membership as having any bearing on whether it is permissible to kill a particular individual, as there are ways of understanding its relevance that are more plausible than McMahan allows.

KEY WORDS: abortion, euthanasia, Jeff McMahan, permissible killing, species membership

Applied ethics is a neglected area of philosophical research. Few, if any, of those graduating from the top philosophy graduate programs in any given year list "applied ethics" as their area of specialization. One reason for this is that research in applied ethics is often thought to lack the kind of subtlety, insight, and rigor that is characteristic of the best work in other areas of philosophy.

One of the many virtues of Jeff McMahan's magisterial new study, *The Ethics of Killing: Problems at the Margins of Life*, is that it conclusively shows what a philosophically rich harvest a talented philosopher working in applied ethics may well hope to reap. Probing, provocative, nuanced, and searchingly honest, McMahan's study is an exemplar of analytic moral philosophy at its very best.

At just over 500 pages, this is neither a quick nor easy read. Each of its five chapters are sufficiently detailed to have served as the core of independent monographs on psycho-physical personal identity and the basis of egoistic concern, the badness of death, the wrongness of killing, abortion, and euthanasia. To a certain extent, this is a weakness of the

A review essay of Jeff McMahan. *The Ethics of Killing*: *Problems at the Margins of Life* (Oxford: Oxford University Press, 2002).

The Journal of Ethics (2008) 12:57–80
DOI 10.1007/s10892-007-9021-8

book, as McMahan's extensive discussion and criticism in each chapter of the views of others tend to obscure the continuity of argument across the different chapters. To be fair, McMahan does his best to remind readers at various points of the central claims for which he has been arguing, but amidst the pages of dense argument and examples, these signposts are not wholly effective.

I say that the density of each of McMahan's chapters is only a weakness of the book to a certain extent because, when one abstracts from the specific arguments of each chapter to reflect on the book as a whole, it is apparent that it is in fact tightly woven together. Conclusions about personal identity and basis of egoistic concern are subsequently deployed in the account of the badness of death, and the findings of both are put to use to argue for the third chapter's pivotal account of the wrongness of killing. The resources of this account are then deployed in the last two chapters to explore the best case that can be made for the permissibility, in certain circumstances, of abortion (Chapter Four) and euthanasia (Chapter Five).

In what follows, I will first sketch the central conclusions of each of McMahan's chapters, in order to bring into relief the ways in which the discussions of the different chapters build upon one another. This is far too rich and detailed a book to do it justice in a review; the discussion of the central line of argument that unites the book will, therefore, of necessity leave out much that is of interest. Having outlined McMahan's central claims, the last part of this review will turn to some doubts, both methodological and substantive, mainly concerning McMahan's theoretical account of the wrongness of killing.

I

The thought that frames *The Ethics of Killing* is that though killing is normally wrong, it is not, intuitively, always wrong. One does not have to resort to controversial examples to see why this is so; that it is permissible to kill in self-defense and to euthanize pets who are in great pain are widely shared moral convictions that point to the existence of exceptions to the general prohibition on intentional killing. What is not so intuitively clear is what the relevant features are of the cases of permissible killing that account for their status as *permissible* killings. McMahan rightly takes this to be among the most difficult questions in moral theory, one that, in his view, has yet to be satisfactorily answered. It is also a matter of great practical importance. It is not unreasonable to think that a plausible account

of when it is permissible to kill that illumines the justification for the permissibility of killing in the uncontroversial cases ought to contribute to a better understanding of what the best justification might be for the permissibility of controversial kinds of killings, such as abortion and euthanasia.

McMahan's approach to the question of when it is permissible starts with a general theoretical account of why killing is wrong. The resources of this account are then deployed to identify the salient features of cases of permissible killing that account for the normal prohibition against intentionally killing not applying to these cases. Presentation of this general theoretical account is undertaken in the first three chapters of the book. The first develops and defends an account of psycho-physical personal identity that McMahan calls the "embodied mind account". It makes two important claims. First, what makes an individual *some one*, rather than just a *thing*, is the capacity for consciousness. Without consciousness, it makes no sense to speak of an individual who has reason to care how her life goes, and whose interests in continuing to live may be thwarted if killed. Second, an individual is the same individual over time to the extent that there is *physical* and *functional* continuity in that part of the brain in which the capacity for consciousness is realized. It follows from this account that because an organism may come into being before it is capable of supporting consciousness, and persist after it can no longer capable of doing so, that my body existed before I came into being, and may well continue to exist as a living body after I am gone.

What, then, are we to make of intuitions elicited by various thought experiments, made famous by Parfit, that support the view that personal identity supervenes on facts about the degree of psychological connectedness? McMahan's strategy is to relocate the force of these intuitions. Psychological connectedness, he argues, has no bearing on the question of personal identity over time, but it is relevant for understanding the basis for egoistic concern. Most of us have a strong interest in our future, and in the normal case the rich psychological connections between an individual now and an individual in the future have an important role to play in understanding the stake that individuals usually have in how their futures go. Psychological connectedness is not, however, necessary to justify the claim that an individual has reason to care about her future. On McMahan's account, the fact that the physical and functional parts of her brain that support consciousness now will be identical to those that do so 60 years from now is a sufficient reason basis for egoistic

concern, though not the degree of concern that it is rational to have in the normal case where there are also various psychological relations that connect myself now and myself in the future.

The strength of this view is apparent when used to illumine the kind of concern that would be appropriate if one had reason to believe that one may fall victim later in life to Alzheimer's disease. On accounts which tie identity to psychological connectedness, there is no reason to take any kind of egoistic interest in one's far future if one becomes a victim of Alzheimer's disease and it eradicates all psychological connections to one's past. McMahan's account does not disagree with the intuition on which this claim relies—that from one's present point of view, one's future self has the character of a stranger. All that shows, however, is that one has *less* reason to take an interest in one's future self, not that one has no reason, for one's own sake, that, for instance, one not suffer in the future. The same line of reasoning applies, *mutatis mutandis*, to one's relations to the early stages of one's life to which one does not normally have any kind of psychological connection.

McMahan's views concerning the badness of death are informed by his account of the basis of egoistic concern. Following Nagel and others, he takes the badness of death *for the person who has died* to be accounted for by appeal to what she has been *deprived* of by death. The evaluation of the badness of death for the person who has died turns, then, on a counter-factual question: how would the person's life have gone had he or she not died. As McMahan deftly argues, how this counterfactual is to be evaluated is a more difficult matter than one might have reasonably expected it to be at the outset.

The relevant complexity turns on the difference between two different accounts of the badness of death, the *Life Comparative* account, and the *Time-Relative Interest* account. The Life Comparative account evaluates the badness of death for the person who has died by considering the goods that would have been realized in her life had she not died. The difference between this and the Time Relative Interest account is that the later claims that the tragedy of a person's death for the person who has died ought to be evaluated in terms of the goods *that the person had prudential reason to care about* at the time of death. The weaker the degree of psychological unity between the person at the time of death and the person at the time the good in question would have been realized in her life, the less important that good is for evaluating the extent to which the person's death was a tragedy *for her*. The Time Relative Interest account, then,

revises the Life Comparative account in light of McMahan's claims concerning personal identity and the basis of egoistic concern.

McMahan's reason for taking the Time Relative Interest account to be superior to the Life Comparative account is that the Life Comparative account has a deeply counter-intuitive implication. It suggests that, in comparing the tragedy of the death of a newborn and the death of a young adult, the death of the newborn must be the greater tragedy (for the person who has died) because the infant has been cheated of more of the goods that life has to offer than the young adult. McMahan, rightly I think, takes this to be a very counter-intuitive conclusion, one that the Time Relative Interest Account avoids. It avoids it because, though it is true that the infant has been deprived of more goods, it is also true that the psychological relations that unify the infant with the later stages of her life are quite weak; at the time of death, then, the infant does not have particularly strong reasons to care about many of the goods that she has been deprived of by death. The same cannot be said of the young adult.

The Time Relative Interest account plays a central role in Chapter Three of *The Ethics of Killing*, in which McMahan argues for his theoretical account of the wrongness of killing. His initial suggestion is that what is fundamentally wrong with killing—here he has in mind the killing of all kinds of animals, including human beings—is that it frustrates the time-relative interests in continuing to live of the victim. To a certain extent, this fits with certain intuitions about killing. The suggestion that the killing of animals is not morally objectionable is morally perverse, but it is plausible to think that that the killing of a person is a much more serious wrong than the killing of an animal, just as the killing of a mouse is less seriously wrong than killing a dolphin or a chimp. This is so for two reasons, both having to do with the psychological capacities of most non-human animals. First, non-human animals do not have the same range of goods available to them; some goods require complex reasoning and planning abilities that non-human animals do not (to the best of our knowledge) possess. Second, the psychological capacities of non-human animals do not allow their lives to be as psychologically unified as the lives of humans can, in principle, be; certain goods, such as living a life that has a certain character, requires the possibility of that kind of unity.

The wrongness of all killings cannot, however, be adequately accounted for solely by appeal to the victim's frustrated time relative interests. No one, for instance, intuitively believes that the killing of a very old person is less morally objectionable than the killing of a

middle-aged person. But that is exactly the conclusion one is led to if one looks only to the interests frustrated by death to account for the wrongness of killing. It is in response to this problem that McMahan introduces the Two-Tier account of the wrongness of killing. This account distinguishes between *persons* and animals. Killing a person, on this account, is normally wrong because it constitutes a failure to respect the victim's status as a person. Though killing a person need not always mark a failure to treat a person with the appropriate respect for her status, the Two-Tier account claims that the seriousness of wrongfully killing a person does not vary with the strength of her time-relative interests in continuing to live. The wrongness of killing a person, then, has to do with failing to respect the person's status, while the wrongness of killing an animal who is not a person has to do with the frustration of the animal's time-relative interests in continuing to live.

This as yet incomplete statement of the Two-Tier account suggests a view that many will find intuitively appealing, that we should be Kantian in our moral thinking about humans, and utilitarian in our thinking about non-human animals. McMahan's division of the moral norms governing killing into the "morality of respect" and the "morality of interests" is certainly not immediately at odds with such a reading. Only humanity, one might say, has dignity, which is of incomparable worth. For this reason, it is impermissible to kill one human being in order to serve the interests of several others. Everything else in nature, including the lives of non-human animals, has price. The life of a non-human animal, then, can be sacrificed in order to bring about some important good or benefit for others. This line of thought certainly makes sense of certain strong and widely shared intuitions; it is never, for example, thought to be justifiable to conduct medical experiments on human beings, while medical experimentation on animals can, in principle (though perhaps not easily as many think it is), be justified by appealing to the interests that will be served by the experimentation.

The problem with this kind of account, which sharply distinguishes it from the Two-Tier account, is that it unquestioningly assumes that all human beings are necessarily persons simply in virtue of being human beings. McMahan argues that the assumption ought to be questioned, and that, under scrutiny, it proves to be untenable. Arguing for a position he calls "convergent assimilation", he advances the view that if an individual is a member of the class of *persons*, it must be in virtue of certain *intrinsic* properties of that

individual. In particular, if an individual counts as a person, rather than just a kind of animal, it is in virtue of that individual having developed certain higher and more complex psychological capacities (of a kind associated with most adult human beings). Many human beings have the requisite higher psychological capacities, but not all do, e.g., cognitively impaired human beings and infants. That is, the class of human beings is *not* co-extensive with the class of persons; the wrongness of killing some human beings ought, therefore, to be explicated, not in terms of a failure of respect, but by appeal to the time relative interests frustrated by the killing. This conclusion has startling implications. If the wrongful killing of a cognitively impaired human being, or infant, is thought to be seriously morally objectionable, consistency requires that we morally assess the killing of a non-human animal *with identical psychological capacities* by the same criteria. Though many tacitly assume it to be true, the belief that killing of an animal is always a less serious wrong than the killing of a human being is one that proves to be just false. Certain cases of killing human beings, such as infanticide may be less morally objectionable than most are intuitively inclined to believe, and the killing of certain animals dramatically more morally objectionable than is commonly believed.

The Two-Tier account of the wrongness of killing is skillfully deployed in the final two chapters of the book, which argue, in turn, for the moral permissibility of abortion and euthanasia. McMahan appeals to the embodied mind account of identity to quickly set aside as unproblematic early term abortions. An early term abortion is one that takes place before the structures of the fetal brain have sufficiently developed to support the capacity for consciousness. Without even the capacity for consciousness, there is no *someone* who can intelligibly be thought of as being killed by an abortion procedure. An early term abortion, then, involves *destroying something*, but not *killing anyone*. The controversial cases, those that can intelligibly be thought of as killing someone, are those that take place after the capacity for consciousness has developed. It is here that McMahan's account of the wrongness of killing becomes crucial for understanding the permissibility of abortion. The psychological capacities of a late-term fetus cannot be plausibly defended as sufficiently developed to require assessing the justifiability of a late-term abortion by appeal to the norms of the morality of respect. How wrong the killing of a late-term fetus turns out to be, if it is wrong at all, therefore turns on the strength of the fetus's time-relative interests

in continuing to live that would be frustrated by an abortion procedure. These interests are, unsurprisingly, quite weak, as there is little to no psychological continuity that unifies one's life as a fetus with the later stages of one's life. If the interests a woman has in having an abortion, then, are sufficiently serious, it is most likely that they will outweigh the interests of the fetus in continuing to live. This conclusion seems to follow even for very late-term abortions. The fetus's time relative interest in continuing to live strengthens as it reaches the end of the gestation cycle. But there is no reason, in principle, that a woman's interests in having even a late-term abortion might not outweigh the interests of the fetus's interest in continuing to live. Abortion after the fetus has developed the capacity for consciousness, then, proves harder to justify, but is arguably justifiable on the basis of a larger set of considerations than is usually acknowledged.

This implication of McMahan's argument is relatively innocuous in comparison to what it tells us about our current understanding of how it is permissible to treat infants. As infants lack the requisite psychological capacities to command respect, so the wrongness of killing an infant must be evaluated with respect to her time relative interests. The time relative interests of an infant in being killed are not, however, significantly stronger than those of a late-term fetus. To the extent, then, that it is harder to justify infanticide than abortion, it is because the interests that tell in favor of infanticide can be satisfied in ways that do not involve killing. Infanticide may not be easily justified, but there is no principled distinction between the permissibility of abortion and the permissibility of infanticide.

What, according to McMahan, constitutes an adequate justification for infanticide constitutes what is undoubtedly the most controversial discussion in *The Ethics of Killing*. The morality of respect, according to the Two-Tier account, is non-consequentialist, or Kantian, in character. It prohibits, for instance, the intentional killing of those who fall within its domain, regardless of the strength of their time-relative interests in continuing to live. The morality of interests, by contrast, is broadly, though not officially, utilitarian. Two respects in which this is so are particularly important. First, the morality of interests allows trade-offs between the interests of those who fall within its scope of a kind that the morality of respect does not permit, such as killing one in order to save many from death. Second, distinctions such as that of doing and allowing, which are

relevant for reasoning about what the morality of respect permits, have no role to play in the morality of interests. As infants fall in the domain of the morality of interests, it is morally permissible to intentionally kill a healthy orphaned infant (who has no living relatives) in order to harvest her organs for the benefit of 10 other children who will die without organ transplants.[1]

The implication of this line of argument for clarifying issues concerning the permissibility of medical experimentation on sentient beings is subtler. McMahan wants to allow for permissibility of medical experimentation on non-human animals, provided the benefit of doing so is sufficiently compelling, and there are no reasonable alternatives. Structurally, this argument is no different than that deployed in the argument concerning permissible infanticide. It does not follow, however, that it is morally permissible to substitute a human infant for an animal in a medical experiment that justifiably employs animals as test subjects. This is not, however, because there is a definite *in principle* prohibition on conducting medical experiments on infants. Rather, there are two main reasons why it is very difficult to justify using infants as test subjects in medical experiments. First, both the amount of good that a normal human infant may reasonably hope to enjoy is much greater than that which lies in prospect for most animals. Therefore, the benefit that stands to be secured by experimenting on an infant has to be dramatically greater than that which could justify using many kinds of non-human animals as test subjects. Second, if there is reason to think that using an infant as a test subject is likely to have long-term traumatic consequences for it that will continue to impair its life even after it has become a person, doing so could be ruled out on the grounds that it would be a violation of the requirements of respect for persons. Arguably, McMahan's position on this matter is at odds with the conviction of some that using an infant as an experimental subject is strictly morally impermissible; to think that the consequences for an infant's life of doing so have any bearing on its permissibility is to seriously misunderstanding the basis of the prohibition. Still, his position on this point is much more in line with commonsense convictions than his conclusions about infanticide.

The conviction that whether a human being is an infant or an adult has no bearing on the permissibility of harvesting her organs in order to benefit a great many others is widespread and quite

[1] McMahan, *The Ethics of Killing*, p. 360.

powerful. One promising line of defense for such a conviction appeals to the *potential* of a healthy human infant; unlike a non-human animal, a human infant has the potential, with suitable care and a bit of luck, to grow into a full-fledged person.

This line of argument fails, according to McMahan, for several reasons, of which I will only mention two. First, it is not true that all human infants have the potential to develop into persons; cognitively impaired human beings do not, for instance, have such a potential. That members of the species *normally* develop into persons is neither here nor there; the morality of killing or harming an individual is to be assessed on the basis of the intrinsic properties of the individual, not her relational properties, e.g., membership in a particular species. Second, though it may be true that, for instance, experimenting on a infant may prevent that infant from ever developing into a person, the infant only has a very weak time-relative interest in becoming a person. McMahan illustrates this point with a nice example:

> Imagine the prospect of becoming like a god. Imagine the possibility of becoming vastly more intelligent and developing a vastly richer and deeper range of emotions, including emotions of which one cannot now form any conception...Even if the transformation would be identity-preserving and would lead to a state that would be clearly superior to one's present state, it would be too much like becoming someone else—and, of course, losing oneself in the process—to be very desirable from an egoistic point of view.[2]

The infant may have, in some sense, the potential to become a person, but the strength of the infant's interest in becoming that person is analogous to the strength of my interest in becoming godlike. Because the interest is weak, it is easy to see how sufficiently strong competing interests, such as the interests in continuing to live of six other infants in need of organ transplants, might well outweigh it.

As disturbing as the implications of this conclusion are, the fact that McMahan acknowledges his own discomfort with them, and that he explores at great length possible flaws in his line of reasoning, is a testament to his commitment to philosophical excellence. A lesser philosopher would not have drawn attention to the more controversial implications of his account of the wrongness of killing that some

[2] McMahan, *The Ethics of Killing*, p. 322. Though I believe this example nicely illustrates the point, McMahan offers it for a slightly different purpose, concerning the relevance of potential as a source of interests, rather than its relevance for questions of moral status. He offers a different set of examples to argue that potential is not a source of moral status, which will be taken up in the second part of this discussion.

will no doubt treat as a *reductio* of his whole approach to the subject. To do so, however, would be a mistake. Hiding from McMahan's arguments will not change the fact that they present a powerful rational challenge to certain moral convictions concerning how it is morally permissible to treat infants.

The arguments that map the terrain of permissible euthanasia mirror those concerning abortion. Because a human being is not necessarily a person, it is possible for an individual human being to cease to be a person, as the individual's psychological capacities deteriorate. If the deterioration continues to be the point that the individual loses even the capacity for consciousness, the only remaining moral question will be whether or not to keep the organism alive. There may be reasons to do so, but they will not be reasons that concern the interests of the individual, as that individual has ceased to exist. That leaves two hard cases. The first concerns the permissibility of persons ending their lives, a question concerning the requirements of the morality of respect. The second concerns the permissibility of denying life-supporting aid to a conscious non-person when doing so is a matter of complying with wishes expressed while she was still a person. Here, the morality of respect and the morality of interests may support different conclusions concerning what is morally permissible.

Why might euthanasia or suicide be thought to be incompatible with respect for the dignity of persons? Kant's discussion of suicide in the *Groundwork of the Metaphysics of Morals* (*Akadime*, 4:429) suggests an argument of this kind, one elaborated in detail by David Velleman, whose arguments concerning the impermissibility of suicide McMahan discusses at length, and ultimately rejects.[3] At the heart of the issue is the question of whether or not it is morally permissible to end one's life if what life has to offer is a life that is in fact a *life not worth living*. What a person's prospects would have to be like to have a life not worth living might, in particular cases, be a matter of some controversy, but I do not think it is implausible to accept that there are such cases. McMahan offers several arguments that favor rejecting Velleman's Kantian argument against the permissibility of suicide. I will not review these arguments in detail, as I share McMahan's sense that the intuitive plausibility of the thought that respecting one's dignity as a person requires that one persist through suffering that will only end at death, is, to say the

[3] David Velleman, "A Right to Self-Termination?" *Ethics* 109 (1999), pp. 606–628.

least, opaque. As he rightly says, the suggestion that ending one's life in order to avoid facing an ordeal at the end of which one will be faced with the task of rebuilding one's life in a way that "will redeem one's previous errors or misfortunes...is certainly to treat oneself as a paltry and disposable thing, unworthy of serious effort or commitment."[4] There is no reason, however, to think that this kind of case generalizes to all cases, especially those in which one's future holds, *in principle*, no prospect of utilizing one's capacity for rational agency in actively guiding one's life in light of self-chosen goals and commitments.[5]

Euthanasia may be compatible with the kind of respect that is owed persons (and that persons owe themselves), but what about cases of advanced directives, particularly those in which a person requests that her life be ended if she ceases to be a person? Here McMahan specifically has in mind the case of a patient, faced with the prospect of dementia, who makes it clear, when her rational capacities are in good order, that she has no desire to continue to live in a state of dementia. Once her psychological capacities have been significantly eroded, however, she expresses a strong preference to continue to live. McMahan argues that the advanced directive ought to be respected, for two reasons. First, in evaluating the force of the individual's later wish to stay alive, we should look to the strength of the individual's time-relative interests in continuing to live; in this kind of case, they will be comparatively weak, a consequence of the dementia eroding the psychological connections that unify her life. And as she is no longer a person, her expression of a desire to continue to live lacks the moral authority of the expressed wishes of a rationally self-governing individual.

Though these interests provide some reason, for the sake of the individual as she is now, to go on with life, McMahan argues that they are not sufficient to trump her expressed will, while still a person, for her life to be put to an end when she has slipped into a state of dementia. This is because,

> ...the earlier part of the life is overwhelmingly the dominant part, its good should have priority. The earlier part was, in itself, a reasonably full and complete life with its own deep prudential unity. It was the life of the individual in her higher

[4] McMahan, *The Ethics of Killing*, p. 484.

[5] This is a position that Kant, at least as I read the *Groundwork*, need not disagree with. The case he cites in discussing the prohibition against suicide is of a person who wishes to escape a "trying condition." A state of intense and unrelenting suffering is surely something different in kind than a "trying condition."

state, when she was a rational and autonomous person. Its good—which is the good of the earlier higher self—is therefore more significant than the good of the shallow and necessarily rather brief period of dementia hanging at the end of life.[6]

Nothing in the Two-Tier account of the wrongness of killing, then, provides a reason to think that euthanasia is a morally impermissible form of killing.

II

Many of the more startling conclusions of *The Ethics of Killing* can be traced to a general methodological commitment to the view that the way in which it is justifiable to treat an individual turns on the specific properties of that individual, and in particular, facts about that individual's psychological capacities. From this it follows that group membership is morally irrelevant. The fact, for instance, that members of a particular group, or species, normally, e.g., have certain psychological capacities, has no bearing on what the psychological capacities of a specific individual happen to be.

Though I do not doubt that facts about an individual's psychological capacities can be morally relevant for thinking about our obligations to, in particular, the cognitively impaired, I have reservations about McMahan's deployment of psychological capacities as having a bearing on questions concerning the permissibility of killing. In what follows, I will first consider a general doubt I have about whether intuitions about the assessment of the tragedy of death at different points in a person's life really lend support to the Time Relative Interest account. Second, I will consider a way of characterizing the relevance of species membership for moral reasoning that is not, I believe, vulnerable to the reasons McMahan offers for it not being relevant. Finally, I will offer some *very* tentative thoughts about how one might defend the claim that the wrongness of killing *any* human being is best characterized as a failure of respect.

One reason for thinking that psychological facts have the kind of ethical relevance that McMahan takes them to have has to do with the role facts about psychological connectedness play in vindicating the Time Relative Interest Account of the badness of death over the rival Life Comparative Account. A key problem with the latter

[6] McMahan, *The Ethics of Killing*, pp. 502–503.

account is that it appears to yield the wholly implausible conclusion that the death of an infant is worse for the infant than the death of a middle-aged person for that person because the child has been deprived of more of the goods that life has to offer. This conclusion is neatly averted by the Time Relative Interest account.

Arguably, this is not the decisive objection to the Life Comparative account that McMahan takes it to be. If something like, say, pleasure, is used as the relevant conception of well-being that is deployed by the Life Comparative Account to evaluate what a person has been deprived of by her death, McMahan's objection to it is undoubtedly a strong one. If a different conception of well-being is presupposed, however, the Life Comparative account proves to be surprisingly resilient. For instance, suppose the Life Comparative account is taken to presuppose an account of well-being that, *very* roughly, takes well-being to consist in an individual's success in the pursuit of valuable activities and goals that matter to her.[7] An account of this kind can be deployed for determining how bad an individual's death was for her at the time of her death, and will not be committed to the view that the death of an infant is worse for it than the death of a middle-aged person is for that person. Rather, it will characterize the way in which the infant's death is bad for the infant as being very different from the way the death of the older person is bad for that person. While the death of an older person deprives her of the opportunity to fully realize many of her goals, the infant has been deprived of the opportunity to develop the kinds of valuable goals and relationships success in which is constitutive of individual well-being. The question of whether or not an infant's death is worse than that of an adult in the prime of life proves, on this account, to be ill-posed.

More needs to be said to develop and defend an account of this kind. Its importance lies in the challenge it presents to the relevance of the kinds of psychological facts that the Time Relative Interest account identifies as relevant for assessing the badness of an individual's death for that individual. This is particularly significant as the Time Relative Interest account plays a pivotal role in the Two-Tier account of the wrongness of killing. The rejection of that account, in favor of a more plausible version of the Life Comparative

[7] Here I have in mind the account of well-being advanced in Joseph Raz, *The Morality of Freedom* (Oxford: Oxford University Press, 1986), though I do not intend the approach to evaluating well-being that I have gestured at to be an adequate gloss on Raz's account.

account, need not, however, lead to a revised version of the Two-Tier account whose implications concerning permissible killing differ from those that McMahan identifies.

Whether or not that is so is an issue I set aside, as it seems to me that it is a mistake to accept that the killing of an infant human being is best cashed out in terms of the frustrated interests of the killed. Frustrated interests may play an important role in accounting for the wrongness of killing at least some (though I doubt all) non-human animals. But the wrongful killing of a human being, in my view, is always best characterized as a failure of a human moral agent to treat another human being (even one who is not a moral agent) as she ought to be treated as a matter of respect for her as a human being.

McMahan offers several subtle and challenging arguments as to why this is an untenable position, whose proper assessment requires more attention than I give them in this discussion. In what follows, therefore, I will first briefly sketch what I take to be one of the more important arguments McMahan offers for rejecting species membership as morally relevant for determining the moral norms that govern how an individual is to be treated, and then offer some tentative thoughts on how it might be defended.

Before turning to that discussion, however, it is important to note that by suggesting that species membership may be morally significant, I am in no way suggesting that membership in the human species endows human beings with a *higher* moral status than any other kind of animal. Those who seek to justify the cruel treatment of non-human animals often appeal to species membership. Such appeals, it seems to me, really are no different than the way racists invoke race or sexists appeal to facts about gender. Ordinary attitudes concerning how it is permissible to treat animals are unquestionably in need of serious revision. It does not follow, however, from species membership not justifying what some take it to justify that species membership is morally irrelevant.

McMahan is sensitive to this line of argument, recognizing that though there is no reason to think that not being human excludes other animals from being entitled to any kind of moral consideration, there might be something to be said for the view that species membership is sufficient, though not necessary, for inclusion in the domain of those individuals whose treatment by others is to be regulated the norms of the morality of respect. The underlying

rationale that unites the norms of the morality of respect, one might argue, concerns the importance of individuals being able to lead rationally self-governed lives. However, to fall under the protection of the morality of respect an individual need not *in fact* have the capacities which make rational self-governance possible. It may well be enough that one belongs to a kind of species that *normally* develops the requisite capacities for rational self-governance.

McMahan's most telling line of objection to this line of argument is nicely illustrated by his example of the "Superchimp." As a consequence of gene therapy, the Superchimp, as an adult, "comes to have the cognitive and emotional capacities comparable to those of a 10-year-old human child."[8] Should we say of this chimp, just because chimps *normally* lack the psychological capacities of a 10-year-old human child, that it is appropriate to treat this chimp as we would another chimp with lesser psychological capacities? How could we justify this? What would we *say* to the Superchimp to justify our treating her differently than human children of comparable abilities?

The strength of the example becomes particularly clear when it is reversed. Say that the population of Superchimps grows to the point that they outnumber the kinds of chimps with which we are presently familiar. If the norm for chimps changes, is it the case that the way it is appropriate to treat the remaining non-Superchimps will have changed? Whatever one might be inclined to think about this question, McMahan's point, which I take to be a good one, is that it is implausible to think that the morality of how it is appropriate to treat an individual should be sensitive to changes in mere *statistical* facts.

If claims about what is "normal" for individual members of a species are just statistical generalizations, McMahan has a strong case for rejecting species membership as having any normative import. But is this the right way to understand the significance of claims about a particular species? To assess the strength of the case, it will be helpful to consider an example. Every spring, I sew several rows of seeds in my vegetable garden. Many of them will never germinate and start to transform into fullgrown plants. But some will. What is the proper way to think about the metaphysics of this process? Has one kind of thing, a seed, been transformed into another kind of thing, a plant? Or is it just that the appropriate conditions obtained in my garden (no small miracle) to enable the seed to realize its nature as a

[8] McMahan, *The Ethics of Killing*, p. 147.

fullgrown plant? The later, I believe, is closer to the truth. Being a seed is just a stage in the normal developmental path of a plant. A complete understanding of a particular kind of plant will include understanding that it is the nature of plants of that kind that they start life as seeds, and if the appropriate conditions obtain, they will realize their nature as full-grown plants. For some kinds of plants, those conditions may be rare, so it might be that the normal developmental path of plants of that kind are frequently interrupted early on in the developmental path. But that does not change the fact that the only thing that would be surprising about a particular seed of that kind realizing its nature by growing into a mature plant would be the statistical rarity of that kind of occurrence.[9]

What examples of this kind suggest, I believe, is that claims about species are not statistical generalizations. Rather, what they concern is the essential nature of a living kind, revealing facts about the *normal life-cycle* of that kind of living thing. The use of ''normal'' here is unashamedly normative. Claims about the life-cycle of a particular kind of living thing, or species, are just constitutive of what it is to be a member of that species. Certain events may, of course, prevent a particular individual member of a species from living out the life cycle that is normal for species of that kind. When this happens, some kind of special explanation is called for, one that explains the deviation from the norm. Some of these explanations will tacitly be informed by an understanding of what members of that species need to flourish. Explaining the death of my rose bush as a result of too much water, for instance, presupposes an understanding of how much water a rose bush needs to remain healthy and grow.

These considerations are suggestive of a possible line of response to McMahan's ''Superchimp'' example. It is important to the example that the Superchimp is definitely of the same species as ordinary chimps. This might be so—biological types allow for a great deal of variation in the characteristics of tokens of a particular a type—but it is not obviously so. And in light of the dramatic difference in capacities that distinguish Superchimps

[9] See a fascinating discussion of this issue in Michael Thompson, ''Apprehending Human Form,'' in Anthony O'Hear (ed.), *Modern Moral Philosophy* (Cambridge: Cambridge University Press, 2004), pp. 47–74.

from run of the mill chimps, it may not even be particularly likely. Arguably, it makes more sense to treat the Superchimp as simply a member of a different species. As a matter of evolutionary theory, that has to be a live possibility. Whether it is or not cannot be settled in the absence of a population of Superchimps sufficient to enable the kinds of empirical observations concerning how they characteristically live required to make a judgment as to whether they are or not a distinct species. If the case can be made that the emergence of the Superchimp marks the emergence of a new species, McMahan will certainly be justified in claiming that it would be morally wrong to apply to the Superchimp the moral norms that regulate the appropriate treatment of other chimps. This does not show, however, that the morality of our treatment of an individual ought to be responsive to the essential properties of that particular individual. It only shows that we ought to be guided by the moral norms that regulate how members of that individual's species ought to be treated.

Much of the intuitive pull of McMahan's critique of the "species membership is sufficient for respect" position draws upon the familiar thought that a person is owed respect in virtue of her capacity for rational self-governance. If, as many Kantian moral theorists do, any appeal to an empirical/noumenal distinction is abandoned, the naturalness of taking the appeal to the capacity for rational self-governance to be a claim about an individual's psychological capacities becomes quite apparent. What McMahan has successfully shown is that if an individual's claim to inclusion in the domain of the morality of respect rests on facts about her psychological capacities, certain classes of human being that many are intuitively convinced ought to be included turn out to in fact have no claim to inclusion.

One way to respond to this insight is to do as McMahan does: accept the conclusion, and skillfully explicate the implications of doing so for commonsense moral convictions. Another response, however, is to re-think the relationship between rational self-governance and the moral norms that Kantians draw attention to in discussing what is respect for humanity, in oneself and in others, requires.

An approach to pursuing the later possibility that I think has promise is to treat the Kantian notion of respect as a more general notion having to do with the reasons to respond in particular ways to the value of something that are required by respect for its

value.[10] For instance, one's appreciation of the value of *The Last Supper* might take the form of planning trips to go and admire it, watching documentaries about it, reading scholarly works that deepen one's appreciation of it, worrying about its deterioration due to age, and debating the merits of various proposals to restore it with others who share your passion. Recognition of this value need not express itself in one's attitudes in these ways, though they are certainly rationally appropriate ways of responding it. But not all ways of engaging with something's value are optional. Some of the reasons for engaging with something of value in particular ways are reasons that all capable of making evaluative judgments are required to take account of in their practical deliberations. A person's indifference to *The Last Supper* does not change the fact that she has reason not to ridicule or disparage it (even in her thoughts), to urinate on it, or to attack it with a can of spray paint. These are reasons that are demanded by *respect* for the value of *The Last Supper*; as they apply to a person irrespective of her particular tastes and inclinations, they can be usefully characterized as categorical reasons.

Just as certain ways of responding to the value of *The Last Supper* are required as a matter of respect for its value, there are requirements that are demanded by respect for the value of human life. It seems to me that what respect for the value of a living thing requires will depend on the characteristic lifecycle, or nature, of members of that species.[11] It is at this point that the Kantian thought that respecting the humanity of another has to do with respect for the other as a rational self-governor becomes relevant,

[10] In this paragraph, I follow an approach advanced in Joseph Raz, *Value, Respect, and Attachment* (Cambridge: Cambridge University Press, 2001), pp. 161–164. Raz's claim is that respect for something's value requires that one have certain attitudes towards it, and that one protect and preserve it. There is no requirement, however, to promote that which is of value.

[11] What respect for the value of the life of different human animals requires is a distinct and open question; it is hard to believe that whatever the best account of those demands are, they will require better treatment of non-human animals than they now receive. Note that the separation of the question of what the moral norms are that ought to govern the treatment of non-human animals from those that ought to govern how humans treat one another marks a significant methodological difference between the kind of approach I am suggesting and McMahan's approach. On his approach, what we have is just one set of norms that governs how animals, as opposed to persons, ought to be treated, whether they are human or non-human.

as an insight that draws attention to importance for human beings of living rationally self-governed lives. That is, respect for the value human life requires taking seriously the importance of rational self-governance in human life. The requirements of the morality of respect, on this approach, then, will have to with more than just guidance concerning how one ought to treat an individual once she has developed the capacity for rational self-governance. It will also concern how one ought to treat an individual to ensure that the right conditions obtain in order for her to develop the capacity for rational self-governance. For instance, most accept that there are duties concerning the kind of care owed children that do not apply to adults.[12] That there are such duties, as well as the specific content of such duties, makes good sense if one accepts that morality of respect is responsive, not just to the particular properties of an individual, but to the characteristic life-cycle of the species to which that individual belongs. Similarly, there will be requirements concerning what respect for the value of human life requires as an individual's capacity for rational self-governance deteriorates.[13]

Part of the appeal of this kind of approach lies in its potential to make sense of certain intuitive moral convictions with which McMahan's view is at odds. For example, there is, according to McMahan, no more reason to administer genetic therapy to a fetus with cerebral defects, at least for its own sake, which will correct those defects than there is to administer the therapy to a

[12] McMahan treats young children who are not infants as falling in an indeterminate zone between those who fall within the scope of the morality of interests and those who fall within the morality of respect. McMahan acknowledges that it is not obvious what to say what morality requires and permits with respect to those who fall into this interdeterminate zone, but takes this to be a problem for all moral theories, not just his. Whether he is correct on this point is an important question, which cannot be pursued here; but, see the discussion of this point in Tim Mulgan, "Critical Notice of *The Ethics of Killing*," *Canadian Journal of Philosophy* 34 (2004), pp. 447–460.

[13] On the approach I am recommending, the defense of the moral permissibility of abortion will have to proceed along the lines of the argument advanced the Judith Thomson, "A Defense of Abortion," *Philosophy and Public Affairs*, 1 (1971), pp. 47–66. McMahan presents a detailed critique of this position, the assessment of which I cannot pursue here.

dog. McMahan flags this as a deeply counter-intuitive, though indisputable, substantive implication of his theory of the wrongness of killing.[14] It follows from the fact that the fetus's time relative interest in developing the cognitive capacities associated with personhood is no stronger than that of a dog (both of which are quite weak). The approach I am suggesting here has the resources to avoid this conclusion. There is reason to administer the genetic therapy to a human fetus, while there is no reason to administer it to a dog, because the therapy will allow the human fetus to develop the capacities that are required to live the kind of life that is appropriate for individual of that species.[15]

An implication of this line of argument is that what respect for the value of human life requires by way of the care and treatment of infants differs from what it requires be done for cognitively impaired infants. Infants who are not cognitively impaired, if suitably cared for and nurtured, are capable of growing into full-fledged adults, capable, if not always willing, to rationally self-govern themselves, and the kind of care and nurturing that morality demands we provide them I take to be responsive to this fact. What morality demands by way of respect for the humanity of the seriously cognitively impaired will be different, to the extent that there is no reason to believe that anything can be done to allow them to develop the cognitive abilities of adults who are not cognitively impaired. Determining what they are entitled to by way of care and consideration is an important, and insufficiently discussed, question, one to which to which facts about psychological capacities are certainly relevant. Their relevance has nothing, however, to do with whether or not they fall into the domain of the morality of interests or the morality of respect. Rather, they are relevant because they help fix the kinds of human goods that an individual, given her capacities, has the potential to benefit from,

[14] McMahan, *The Ethics of Killing*, p. 319.

[15] One natural line of argument at this point is that the cognitively impaired are simply a different species than those who lack cognitive defects. This strikes me as an implausible suggestion. Cognitive impairment is usually understood as the consequence of something having gone wrong. Hence the cognitively disadvantaged are members of the human species who have not developed the normal capacities for human beings, rather than those who have the normal capacities for members of their species.

and that morality requires we help enable her to appropriately engage with.[16] McMahan's account does not treat the cases of the cognitively impaired infant and the healthy infant as significantly different, but that there is a difference is something I find intuitively plausible.

The line of argument I have been exploring might be thought to not touch the core of McMahan's rationale for distinguishing the morality of respect and the morality of interests. The standard kind of argument deployed by non-consequentialists as to why it is normally wrong to intentionally kill another human being, regardless of the strength of her reasons for continuing to live, is that doing so is to make a decision on behalf of the other concerning how her life is to go that is not yours to make. Killing another, in other words, is an offense against that person's status as a self-governing, or autonomous, being. But if a particular human being not only presently lacks the cognitive capacities for rational self-governance, but also lacks the *potential* to develop such capacities, why think that the standard non-consequentialist explanation of the prohibition against killing has a bearing on the justifiability of killing her?

The issue here, as I see it, turns on a difficult and deep question concerning the structure of moral reasoning, one that is roughly structurally similar to the point over which act and rule consequentialists disagree. The act-consequentialist thinks that if the promotion of well-being is what matters, then what we should be attentive to is what course of action will, on any particular occasion, best advance that goal. McMahan's strategy in arguing for the Two-Tier view is one that I take to be relevantly analogous, holding that a prohibition against killing a human being that appeals to the importance of respecting autonomy is only relevant to the assessment of the permissibility of killing an individual when she in fact has the relevant psychological capacities that constitute a capacity for rational self-governance.

According to rule consequentialists, where the act consequentialist has gone wrong is that they do not take the role of norms, or principles, in moral reasoning seriously enough. In particular, they think that what is appealed to in justification of a principle, or system of principles, is not the rationale that is relevant for justifying all the

[16] I do not think that much is to be learnt from reflection on the case of Anencephalic infants. As McMahan himself points out, there are no firm commonsense convictions regarding what is morally permissible with respect to these infants (McMahan, *The Ethics of Killing*, p. 208).

particular acts a particular principle permits or requires. The reason for being guided in one's practical deliberations by a valid principle is that it is a valid principle.

The general interest for the issue at hand of this family dispute among consequentialists is that, it seems to me, non-consequentialists who believe that there are certain requirements that apply to how it is permissible for human beings to treat one another agree with the rule consequentialists concerning the structure of moral reasoning. That is, their position is best interpreted as holding that when we reflect upon the rationales for the moral principles that spell out what we owe one another as a matter of respect for the value of human life, we appeal to values like the importance of being able to lead a rationally self-governed life. The principles, however, are for the general regulation of how humans interact with one another. An individual falls within the domain regulated by these principles, then, simply in virtue of being of human born. Something like this is what I take Kantians to have in mind in portraying humanity as moral community, one bound together by common laws.

This way of thinking about morality is one that at least I find to be very intuitively compelling. If it can be shown to be defensible, it promises to provide a powerful framework for arguing for some of the common moral convictions that McMahan calls into question that even those with very liberal moral views will be reluctant to abandon. It also, I believe, identifies an important range of questions concerning the ethical treatment of animals. Just as there are norms that regulating how human beings treat one another that flow from the requirements of respect for the value of human life, there will be norms concerning what respect for the value of different species of non-human life require. Whatever these requirements turn out to be, they will no doubt turn out to require much more of us than most are currently prepared to accept.

There is much in *The Ethics of Killing* that is both controversial and deeply counter-intuitive. Its brilliance lies in the pages of rigorous and thoughtful argument that McMahan offers in support of his conclusions. Together, they constitute both a powerful defense of certain moral convictions that at least some will share, and an important philosophical challenge to the plausibility of other convictions that even those with very liberal moral views will be reluctant to accept as being mistaken. Assessing whether this

challenge can, even in part, be met requires engaging in the kind of systematic, detailed investigation of the sorts of questions pressed by McMahan concerning the foundations of morality that meets the very high standard he has set.[17]

Department of Philosophy
Queen's University
Kingston, ON, Canada K7L 3N6,
e-mail: kumar@post.queensu.ca

[17] For helpful discussion of many of the issues discussed here, I am grateful to the students in my graduate ethics seminar in the fall of 2004. Also, thanks to Jeff McMahan for helpful clarification of his views at crucial points.