

MCMAHAN ON THE WITHDRAWAL OF LIFE-PROLONGING AID

INGMAR PERSSON

The University of Gothenburg

JULIAN SAVULESCU

The University of Oxford

1. Introduction

In his landmark book, *The Ethics of Killing*,¹ Jeff McMahan provides the most thorough, comprehensive and imaginative analysis of killing “at the margins of life”, i.e., killing of human beings who are not persons, so far constructed. We find ourselves in agreement with the great majority of McMahan’s arguments, which significantly advance debate in many seemingly intractable areas. His arguments are many and subtle, and the literature covered so vast, that we cannot hope to consider anything but a tiny fraction of his positions in this essay.

We will focus on his account of the distinction between killing and letting die, the KLD-distinction. This distinction crops up in his treatment of both abortion and euthanasia. We are in sympathy with McMahan’s aim to give an account of the KLD-distinction which is morally neutral in the sense that it does not presuppose that killing is harder to morally justify than letting die. But we will suggest that he fails to achieve this aim when he contends that cases of death issuing from the withdrawal of life-prolonging aid that one provides are instances of letting die rather than of causing death. It is our suspicion that he unconsciously classifies these cases as instances of letting die *for the reason* that withdrawal is here permissible.

This sort of misclassification could create the illusion that the claim that the KLD-distinction has moral significance is more plausible than it actually is. Although McMahan does not definitely commit himself to the truth of this claim, he seems to us more respectful of this claim than it deserves (p. 383). Moreover, this claim could in turn perpetuate the standard view that active euthanasia is harder to morally justify than passive euthanasia, though McMahan uses it rather to the opposite effect (p. 461). For our own part, we firmly reject the claim that the KLD-distinction and, hence, the distinction between active and passive euthanasia, are morally significant. We will give some arguments in support of this rejection, but at the heart of our interest lies the analysis of the KLD-distinction.

1. All page numbers in parentheses refer to this book.

2. The Killing/Letting Die Distinction

According to common sense morality, killing human beings is morally worse, or harder to justify morally, than letting them die; so there will be circumstances in which killing human beings will be wrong, but letting them die permissible. For instance, while it would be permissible to let one human die by the roadside in order to have time to save two, it would not be permissible to run over one in order to have time to save two. We will call the claim that the KLD-distinction has this sort of moral significance, *the moral significance claim*. A natural corollary of this claim is the view that active euthanasia is morally worse than passive euthanasia, *the euthanasia corollary*. This is a natural corollary, since it is plausible to define active euthanasia as killing, or causing death, with the end of benefiting the victims, and passive euthanasia as letting die with the end of benefiting the victims. Active euthanasia is widely considered impermissible and is illegal in virtually all countries (the Netherlands and Belgium are two exceptions), whereas forms of passive euthanasia are generally considered permissible and are legal in many countries.

As McMahan is well aware, it is important not to allow the moral significance claim to confuse the issue of what the KLD-distinction consists in or how it should be analysed. For instance, it is important not to assume that some piece of behaviour must be a killing of a human being rather than an instance of letting such a being die because it is morally wrong or, conversely, that it must be an instance of letting die rather than a killing because it is permissible. We take it as obvious that, for instance, *killing* is not a normative or evaluative concept because non-moral agents, like earthquakes and predators, can kill. Advocates of the moral significance claim probably like to think that some piece of death-oriented behaviour is, in most circumstances, wrong *because* it is a killing of a human being or permissible *because* it is a letting die of a human being, and they cannot (non-circularly) do this if they hold that something can be a killing of a human being only if it is wrong, and a letting die of such a being only if it is permissible. So, like McMahan, we regard it as a desideratum of an analysis of the KLD-distinction that it is morally neutral. Let us now propose an analysis that obviously meets this desideratum.

Killing is a species of the genus of causing death. Hence, the KLD-distinction is an instance of the more general distinction between causing something (to occur) and letting it occur or be caused. For present purposes, it will have to be enough to explain the general distinction with sufficient precision to make possible a contrast between what it implies for the KLD-distinction with certain other conceptions of that distinction. We suggest that at t you let it be a fact that p (let an event occur or a state obtain) if and only if (a) p will be a fact unless some *change* is made to the way you at t correctly take things then to be, (b) at t you know (or correctly believe) that you *can* (in an 'all-in' sense entailing both ability and opportunity) then prevent p from being a fact by *causing* some change, (c) you choose or decide not to cause any such change at t , and (d), because of (c), p is a fact at t . This account uses as primitive the intuitive notions of a (genuine not merely a so-called Cambridge) 'change', of 'causing' and of 'can'. It implies that, when you let it be a fact that p , you

make a decision with respect to p , such that if things had not been in accordance with that decision, but you had instead made no decision, it would still be a fact that p . In other words, when you let something be the case, things are in this respect as they would have been if you had made no decision which was implemented, if you had had no capacity to decide and behave accordingly.²

As an illustration of how moral judgement can influence the classification of something as a killing or an instance of letting die, McMahan takes the case of a mother who fails to feed her infant, with the result that it dies of starvation. There is doubtless a rather widespread tendency to describe this mother as having killed her infant. McMahan thinks this is a misdescription which is due to one's being influenced by a moral evaluation of the case: since the mother is under an obligation to take care of her child, her omission to feed it is wrong like a killing; so, it *is* a killing (pp. 236–7, 386). We agree with this diagnosis. On the account of letting happen we have just presented, the woman does not kill or cause the death of the baby because the relevant decision she made, the decision not to feed her baby, is such that if it had not been made and implemented, but she instead had made no decision—if, for instance, she had been unconscious or dead at the time—the baby would still have died. In this sense, the mother's decision and its implementation made no difference to the fate of the baby.

Having this goal of a non-moral analysis of the KLD-distinction means that we side with McMahan against, for instance, Judith Jarvis Thomson who contends that “the difference between killing and letting die is not wholly nonmoral”.³ She proposes the following as a necessary condition of letting a patient die: the agent “has a liberty-right to engage in the behavior (action or inaction) that issues in the patient's death” (p. 504). If this is a necessary condition, it follows that the mother does not let her baby die, for she is not at liberty to starve it to death. Our disagreement with McMahan comes from the opposite direction so to speak; it consists in that he is not wholly successful in cleansing the KLD-distinction of moral influence. In the next section, we will argue that he errs by classifying some cases of killing or causing death as instances of letting die, probably because he is unconsciously influenced by the consideration that they are permissible.

3. Withdrawal of Life-Prolonging Aid as Causing Death

The disagreement between us does not concern cases in which life-prolonging treatment is never provided to somebody in need of it. Suppose a person is admitted to hospital with leukaemia in its advanced stages. His symptoms are relieved with simple pain killers. The question arises whether to provide him with aggressive medical treatment in the hope of prolonging his life. Such

2. This implication holds for certain cases of overdetermination in which you decide to cause some change, e.g., to move a finger, but this change would have occurred even if you had not made and implemented the decision (e.g., the finger would have moved as the result of a twitch). It is the content of your decision which rules out that this is a case of letting happen.
3. ‘Physician-Assisted Suicide: Two Moral Arguments’, *Ethics*, 109 (1999), pp. 497–518, pp. 504–5.

treatment will consist of aggressive chemotherapy, followed by intensive care. This would involve drugs to support the patient's heart and artificial ventilation, and other invasive treatment. If this life-prolonging treatment is never started, and the patient dies of his leukaemia sooner than he would have if it had been started, we uncontroversially have a case of letting die and, thus, of passive euthanasia, if the end of this omission is the well-being of the patient.

Our disagreement concerns some cases in which this treatment is started and then stopped or withdrawn. McMahan spells out his analysis of the relationship between withdrawal of life-supporting aid and letting die in this way:

Let us call someone who provides life-supporting aid or protection the *Provider*, and someone who removes or withdraws life-supporting aid or protection the *Remover*. I claim . . . that when a Remover withdraws the source of a person's life support and as a consequence the dependent person dies, he merely let the person die if he is also the Provider of the aid he withdraws. (p. 380)

In an earlier paper, he supplies some further details:

if a person requires or is dependent for survival on further aid from or protection by an agent, and if the person dies because the agent fails to provide further aid or withdraws his own aid either while it is in progress or before it comes operative, and if the agent is not causally responsible for the person's need for aid or protection, then the agent lets the person die.⁴

Thus, suppose that in order to prevent someone, the *Victim*, from dying of some disease, you have hooked him up to your body. If you disconnect yourself, with the result that the Victim dies of the disease, then, on McMahan's view, you have merely let the Victim die, though you have *actively* freed yourself, *caused* it to be the case that you are no longer connected to this individual, with fatal results. Here what you do—the disconnecting movements—make a difference because if they had not occurred, the Victim would have (stayed connected to you and) lived on. So, contrary to our analysis, McMahan's analysis implies that you may let die, though you cause a change (take away aid) which leads to death.

In the case just mentioned, the Provider of life-supporting aid is identical to the Remover of this aid. This identity is not necessary for the case to be one of letting die, according to McMahan. Suppose that a doctor has disconnected the Provider from the Victim at the Provider's request. Then it could still be a case of letting die, on McMahan's view, because the doctor acts on behalf of the Provider.⁵ But, for the sake of simplicity, we may largely disregard this complication and concentrate on the situation in which the Remover is identical to the Provider.

4. 'Killing, Letting Die, and Withdrawing Aid', *Ethics*, 103 (1993), pp. 250–79, at p. 261.

5. 'Killing, Letting Die, and Withdrawing Aid', pp. 264–5; and *The Ethics of Killing*, p. 382.

In the second quotation above, McMahan also requires that the Provider “is not causally responsible” for the Victim’s “need for aid or protection”. But, as he points out elsewhere, this is not really necessary for there to be an instance of letting die. Imagine that the Provider is causally responsible for the Victim’s deadly disease, perhaps because he has infected him. (Call the person playing this part the *Initiator*.) If he then disconnects himself, he may make it true that he has killed the Victim by infecting him. But McMahan makes it clear that he thinks that the act of disconnecting *in itself* is an instance of letting die (p. 380). In this situation, you would let your earlier act of infecting the Victim develop into a killing, on his view.

According to McMahan, one factor which determines whether a withdrawal of operative life-supporting aid is killing or letting die is whether this aid has become “self-sustaining” or is still “in progress”, “continuing” or “requires more from the agent”. At first blush, it might appear that this factor is purely factual, a matter of whether further action from the Provider is needed to keep the aid operative. But McMahan explicitly makes it clear that aid may not qualify as self-sustaining, although no “further *action* is required” to keep it going. It is enough that it uses “resources that properly belong to the Provider” (p. 381). Thus, he believes that disconnecting the Victim from mechanical life-support, or turning it off, will exemplify letting him die when keeping the machine functioning “draws continuously on the provider’s resources and exacts opportunity costs from him”.⁶

But, to our minds, this reveals that McMahan’s analysis of the KLD-distinction is guided by moral considerations. For “resources that properly belong to the Provider” are surely resources that he can *rightfully* or *permissibly* use. So, we are left with the impression that some cases of withdrawal of aid are classified as instances of letting die by McMahan because withdrawal is here permissible. Withdrawal is permissible for the reason that it concerns things belonging to the one who withdraws them, so that he acts within his rights when he withdraws them.

On the analysis of the KLD-distinction we propose, we have a case of killing or causing death—and so active euthanasia—irrespective of whether the operative aid is in progress or self-sustaining. We do not conclude from this, however, that acts of disconnecting from or turning off life-support machines are therefore wrong, for we are sceptical of the moral significance claim and its euthanasia corollary. We hold it to be absurd to think that it cannot be permissible to withdraw aid that it was permissible never to supply in the first place, although the former is causing death and the latter is uncontroversially letting die.

To bring out one implication of McMahan’s analysis that we find counter-intuitive, suppose the Provider refrains from removing the aid in progress that the Victim receives. He then clearly *lets* the Victim live on, by letting the aid remain where it is. But it is peculiar that the alternative to what is on McMahan’s analysis letting the Victim die consists in *letting* him not die or live on. One would have thought, as our definition in section 2 implies, that letting it be a fact that p must have as an alternative an *action* of causing

6. ‘Killing, Letting Die, and Withdrawing Aid’, p. 266.

something that would have prevented p from being a fact. It cannot be that however one behaves with respect to p , one lets either p or not- p be the case.

McMahan agrees that there would be such an alternative action if the Remover is distinct from (and unauthorised by) the Provider. He affirms that, while the Remover would let the Victim survive by letting the aid be in operation, he would kill the Victim if he removed the aid (given that he is not authorised by the Provider to remove it). But, surely, there is a solid sense in which what the Remover does is the same, irrespective of his relation to the Provider, and the rights he thereby may acquire. This is so, just as there is a solid sense in which what a woman does in abstaining from feeding a baby, with the result that it dies, is the same irrespective of her relation to its mother, whether she is identical to the mother or related to her in some other way that puts her under a duty to take care of the baby. The difference lies in the moral evaluation of these pieces of conduct, the removal and the failure to feed, that the Remover acts permissibly if it is his aid to remove and that the woman behaves impermissibly if she is under a duty to cater for her child. Our conjecture is therefore that in classifying cases of permissible withdrawal of life-prolonging aid as instances of letting die, McMahan is guilty of a mistake he himself warns against, namely, that of being misled by moral considerations, just as are those who regard the mother as killing her baby because she acts wrongly.

McMahan's analysis of the KLD-distinction may make the moral significance claim and its euthanasia corollary seem more plausible than they in reality are. For if it were recognised that cases in which Removers permissibly remove their own aid with the result that the Victims die were cases of killing or causing death and, so, cases of active rather than of passive euthanasia, these moral doctrines would call out for revision. This cover up is common enough. Consider the famous case of Tony Bland (*Airedale NHS Trust v Bland* [1993] 1 All ER 821). Bland was in a persistent vegetative state when Law Lords authorised a removal of his feeding tube which resulted in his death. They ordered this because "it could not improve his condition and was burdensome by way of its invasiveness and indignity". However, they considered this to be a case of letting die and, so, passive euthanasia. But in contrast to omitting to put up another bag of artificial nutrition when it ran out naturally, this withdrawal of treatment is causing a change (from Bland's having a feeding tube to his having none) that leads to death. It is true that in cases like this, we can say, 'We removed the feeding tube and let him die.' But this does not imply, as Philippa Foot seems to think,⁷ that we did *nothing but* to let die. Quite the contrary, it explicitly says that we did something else (removed the feeding tube) before letting die. If, however, the case had been recognised as one in which we cause death, the court would have had to face the problem of explaining what relevantly distinguishes this case from cases of active euthanasia and murder which it would hold to be impermissible and illegal.

7. See 'The Problem of Abortion and the Doctrine of the Double Effect' reprinted in John Martin Fischer and Mark Ravizza (eds.): *Ethics: Problems and Principles* (Harcourt Brace Jovanovich), 1992, p. 64.

By morally charging the KLD-distinction, one could then give the moral significance claim and its euthanasia corollary a greater semblance of credibility, by achieving a more extensive fit between this distinction and the distinction between what is permissible and impermissible. But this more extensive fit is of course bought at the price of rendering the support the KLD-distinction gives to these moral doctrines circular. Furthermore, even if McMahan's extension of the class of letting die were accepted, cases of permissible killing or causing death would still have to be countenanced, alongside cases of wrongfully letting die, such as that of the mother who fails to feed her infant. This is what we will now argue.

Suppose we have started a machine that regularly supplies the dying Victim with a drug which has a good chance of being substantially life-prolonging. But this drug is such that the body thoroughly adapts to it. So, if it is withdrawn the Victim quickly dies of a sort of shock and not of the underlying disease. We now discover that the drug does not prolong the Victim's life: in a few minutes, at time *t*, he will slip into coma and die shortly afterwards, without having woken up. If we now switch off the machine, the Victim will die of shock at *t*. To withdraw the aid in this case, *Shock*, would clearly be to cause the Victim's death⁸ and, so, constitute an example of active euthanasia (if the end is to benefit the Victim). It would be a case of causing death or killing because a new cause of death is produced by the withdrawal.

If other things are equal—in particular, if the Victim's conscious life up until *t* is equally good in both cases (and the brief period of unconsciousness equal to being dead for the same period)—it would be unreasonable to maintain that we are not permitted to withdraw the aid, but must let the Victim die of the underlying disease.⁹ If things instead would be ever so slightly better for the Victim if the aid is withdrawn, and other things are equal, it would likewise be unreasonable to deny that this is what we have most reason to do. But if this is correct, the moral significance claim must be false: it cannot always be harder to morally justify killing than letting die, because to kill is what we do if we remove the aid in *Shock*. For if killing is harder to justify than letting die, killing cannot be permissible in a case in which there is no other difference between killing and letting die and required in a case in which it would make things ever so slightly better.

Another way to argue for this view is by comparing a situation in which we let ourselves kill with a situation in which the active withdrawal of aid leads to death. It seems the former could not be harder to morally justify than the latter which many regard as permissible (whether or not they believe it is

8. That this is a case of killing or causing death does not emerge so clearly from McMahan's analysis as from Frances Kamm's similar analysis of withdrawal of aid; see *Morality, Mortality* (Oxford University Press, 1996), vol. II, esp. p. 29. Note that Kamm concedes that this killing would have "practically the same moral weight as the letting die by actively terminating aid".

9. It is noteworthy that the position McMahan takes up carries the reverse implication, that it would be better, more praiseworthy or imperative to withdraw the aid in *Shock*, for by killing we would have actively benefited the Victim (p. 461). This implication strikes us as counter-intuitive. Surely, the fact that, in a euthanasia context, the withdrawal of aid produces a new cause of death rather than merely enables a pre-existent threat to kill cannot make withdrawal more praiseworthy or imperative, if other conditions are equal.

an instance of letting die). In *Allergy* the only chance of prolonging the life of the Victim is by providing some aid A which consists in starting some machine that regularly supplies him with a certain drug. There is however a small risk that if the Victim receives A he will suffer an allergic reaction and die if A is not quickly removed. We apply A because otherwise the Victim will soon, at t , lapse into unconsciousness and shortly afterwards die, without having regained consciousness. Unfortunately, we find that the Victim gets the allergic reaction which will kill him at t if it A is not immediately removed.

It is clear that if we do not now remove A, we will have killed the Victim by applying it. By letting A remain where it is, we let it be the case that we will have killed the Victim (by applying it). If we choose this alternative, what we now do is to let the Victim die or be killed by the allergic reaction, but by letting this happen, it will be true that a past action of ours (the application of A) will be a killing of him. On the other hand, if we remove A, then, on McMahan's analysis, we will let the Victim die from the original disease. Consequently, this analysis implies that this is a situation in which we let the Victim die whatever we do, in which we cannot possibly avoid letting the Victim die, from one cause or another. But, as we have already remarked, it seems that the logic of the notion of letting happen does not allow that, whatever we do, we let something happen, for example that the Victim dies from some cause or another. As our explication of the KLD-distinction in section 2 implies, a situation in which we let something happen, that the Victim dies from a certain cause, for instance, is necessarily a situation in which we could *act* so that this does not happen.

We have assumed that it is permissible to remove A (irrespective of whether it is a case of letting die), but is it also permissible to let A remain in operation, with the result that the Victim dies of the allergy? This seems difficult to deny, if the consequences for the Victim (and others) would be of equal value whichever alternative we choose. Suppose instead that it would be better for the Victim—even if only marginally—that we do not remove A, and that for this reason he desires it, whereas it makes no difference for others. Then this is surely what we should do. But if so then, contrary to the moral significance claim, the fact that we have then killed the Victim cannot make this alternative harder to justify, for the slightest tip in the balance of reasons in favour of this alternative makes it the one we should do.

If, however, it is permissible to let oneself kill by letting A remain in operation, mustn't it be permissible to kill by putting it into operation? Suppose the drug is not automatically administered, but that we would have actively to administer every individual dose of it. If it is permissible to let the drug continue to be administered, though the Victim would die from it at t , it is difficult to deny that, other things being equal, it would be permissible to administer a last dose of the drug with the result that he dies at t . This is however a close parallel to a standard example of active euthanasia, by lethal injection. On the strength of such reasons, we are inclined to think that the moral significance claim and its euthanasia corollary cannot be sustained. But we will not continue to press this point, since McMahan is not committed to these doctrines.

According to our account, we will not let the Victim die, but cause his death, were we to remove A. We do not say that we will then *kill* the Victim. What will kill the Victim is rather the underlying disease (and what may have caused him to have it). McMahan apparently assumes that when an unauthorised Remover removes aid with the result that the Victim dies, he *kills* the Victim, but this has counterintuitive consequences. Suppose the Remover is identical to the Initiator, the person who is responsible for the threat of death which is let loose by the removal. Then he would kill the Victim *both* by producing this threat of death and by removing the life supporting aid. This seems to charge him with one act of killing too many. Moreover, it would seem to follow, absurdly, that he kills the Victim, by infecting him with the disease, *because* he does something else, removing the aid, that also kills him.

Consequently, not every act of causing, or causally contributing to death, is a killing. We then need to distinguish killings from other instances of causing death. It is most difficult to do this with precision. Let us say, vaguely, that what kills is what we would single out among all causal conditions as 'the cause of death'. This may be either an event, like a disease or other pathology that an organism has, or a more concrete entity, like a person, who causes this event. (As already indicated, we regard it as mistaken to believe, as some do, that only people or morally responsible agents can kill.) On this account, a removal of life-prolonging aid would not be a killing, unless it produces the cause of death, as in Shock. Otherwise, it will cause or enable a pre-existent threat of death to kill.

It bears mentioning that on 'the letting side' there is a similar distinction between letting die and letting be killed (by some cause). If I refrain from stopping you on your way to kill someone, I may let him be killed by you, but I do not let him die. It is only when you have, say, stabbed him that I may let him die, by not stopping the bleeding. Letting a victim die seems to presuppose that the cause of death is already in operation, leaving him dying or struggling for his life.

Admittedly, cases of killing and other cases of causing death shade into each other, and the same goes for cases of letting die and letting be killed. This is because of a fuzziness in the notion of the cause of death. This fuzziness is both diachronic and synchronic. Sometimes, but not always, the cause of a cause of the Victim's death or of what kills him will itself be a cause of his death or be something that kills him. For instance, if a snake's poison causes the death of the Victim, or kills him, in a few minutes, not only the poison but the snake's bite, and the snake itself, may be said to have killed or be the cause of the Victim's death. In contrast, if you light a bushfire that several days later kills the Victim, it goes against the grain to say that you killed him by lighting the fire, though you caused it to kill him.

Similarly, if the issue is considered synchronically instead of diachronically. If you deal a haemophiliac a light blow, it is unclear if the cause of his death is the blow or the haemophilia. Both are necessary conditions in the circumstances. If you deal a normal human being a severe blow with a hammer, we do not hesitate to appoint the blow the cause of death, though of course in this case, too, the victim's constitution was among the necessary causal conditions.

However, we cannot see that these unclaritys matter in the present context, since we believe that nothing of moral importance hinges on the distinction between killing and other cases of causing death or between letting die and be killed.

McMahan puts forward an objection to the kind of analysis of the KLD-distinction we advance:

If we insist that *doing* something to a person that results immediately in his death must count as a killing, we will be committed to the absurd conclusion that one would kill a person if one were to actively resist his efforts to save himself at one's own expense (for example, if one were to resist a person's attempt to use one's body as a shield against a lethal threat). (p. 384)

But, in our opinion, this is not absurd if we keep in mind that killing can be permissible. It may be permissible to kill somebody in self-defence, though when what one does is to "resist a person's attempt to use one's body as a shield against a lethal threat", one rather causes or enables this threat to kill than kills oneself. Imagine that the shielded person successfully resists your attempt to shield him. He firmly holds you in between himself and a gunman, so that you are killed by the bullets. Then McMahan would probably agree that the shielded person causes your death, just as he would if he had grabbed you and pulled you into the line of fire. But it might seem curious that if there is a struggle between a shield and the shielded, there is a letting be killed if the former is victorious, but a causing to be killed if the latter is. This fails to capture an evident symmetry between the two parties of the struggle.

McMahan also mentions his case of the Aborted Rescue in connection with the quoted remark. This is a case of a rescuer who pushes off a drowning man when this man panics and threatens to drown the rescuer as well. In his discussion of this case, McMahan claims: "It seems absurd to suppose that his initial efforts to effect the rescue could make him a killer" (p. 380). We do not see why, if this killing is as permissible as letting the man drown in the first place—which certainly would have been permissible had one known in advance that the drowning man would attempt to kill one if one tried to save him.

Let us summarise our main claims with respect to the removal of life-supporting aid. When you remove life-supporting aid, with the result that the Victim dies of an underlying disease, you cause him to die of this disease rather than let this occur, irrespective of whether the aid is yours to remove. For there is a change you make occur such that, if it had not occurred, the Victim would not have died when he did. It is however the disease rather than you which kills. If you act with the end of benefiting the Victim then, since you cause death, it will be active rather than passive euthanasia. Since a removal of life-supporting aid can be permissible, this makes the moral significance claim and its euthanasia corollary hard to sustain. Moreover, we have also suggested that killing can be permissible in the same circumstances in which causing death by withdrawal of aid is permissible.

4. Merely Extractive Abortions

As regards abortion, McMahan's analysis of the KLD-distinction implies that abortions by hysterectomy or hysterotomy that remove the foetus whole—"merely extractive abortions"—can be instances of letting the foetus die "of its own inherent vulnerability" (p. 384). In contrast, we claim that these abortions instead *cause* the foetus to die "of its own inherent vulnerability".

This may be seen most clearly if we consider behaviour related to *one's own death* because the KLD-distinction is not here associated with any moral difference: according to common sense morality, it is not considered to be morally worse to kill oneself than to let oneself die. So, there is no risk that the classificatory issue be muddled by moralising. The case of extractive abortion may be compared to a case in which you walk out of a shelter in which you keep snugly warm into a raging blizzard and succumb to the cold. In the latter case, you die as a result of your "own inherent vulnerability", like the extracted foetus. You 'withdraw' life-preserving aid which you yourself provide and which is "in progress" because you need to feed the fire with wood to keep it heating your shelter. Nevertheless, you clearly do not let yourself die. Through your action, you cause yourself to be exposed to the cold. In contrast, you would have let yourself die of cold if you had remained in the shelter, but had refrained from keeping the fire going by not feeding it with wood.

It might be objected that there is a difference between you and the foetus. When you walk out in the cold, like Captain Oates in Scott's expedition to the South Pole, the cold rather than your own inherent vulnerability is commonly thought to kill you. In contrast, it might be proposed, the foetus is killed by its own inherent vulnerability rather than by the extraction from the womb. This is because the foetus's vulnerability, like the haemophiliac's, is greater than normal. In reply, we claim that there is no need to pick out norms which make Captain Oates's vulnerability normal and the foetus's higher than normal. Captain Oates's vulnerability to cold is normal only if it is compared to other adult human beings—not, for instance, if the comparison includes animals like the Imperial Penguin. Likewise, the foetus's vulnerability might be normal if it is compared to other foetuses. The main consideration is, however, that whatever we single out as the cause of death, or that which kills, death is *caused* in both the Oates case and the extracted foetus case. For even if merely extractive abortions do not kill the foetus, as other forms of abortion do, but cause something else to kill it, this is a distinction in which it would not be safe to invest any moral significance. In any event, we cannot rely on the moral significance claim to morally justify merely extractive abortions, since they are not instances of letting die.

McMahan agrees that there is no moral difference between merely extractive abortions and other kinds of abortion. Since he thinks merely extractive abortions are instances of letting die, while other abortions are killings, he is left with the problem of explaining why the KLD-distinction does not have its customary moral implications in the case of abortion. He suggests that this may be because "the distinction between killing and letting die derives its

significance from the requirement of respect for persons” (p. 389). Plainly, foetuses are not persons in the sense of beings with higher mental faculties like self-consciousness. But this is neither a necessary nor a sufficient condition for the KLD-distinction to have moral significance. It is not sufficient because, as we have just seen, the KLD-distinction does not have moral significance as regards one’s own death. And it is not necessary because, if the distinction has moral significance, it surely has so in the case of new-borns, even if they are prematurely born. Like foetuses, they are not persons, but at least according to common sense morality, it is worse to kill new-borns than to let them die.

There are other counterintuitive implications of McMahan’s analysis of the KLD-distinction, which he acknowledges. It implies that a person who removes fish he has bred from his aquarium (thus removing life-support in the form of water which he has provided) is letting the fish die rather than killing them or causing their death (pp. 387–8). By contrast, if he were to take fish out of the sea, he would be killing them or causing their death. But it seems peculiar that the source from which you take the fish could make the difference between killing them and letting them die. Suppose that one fish from the sea has slipped into the aquarium. McMahan would then have to say that, when the owner removes all the fish from his aquarium, he lets all of them die, except for one which he kills!

In this and the foregoing section, we have argued, in opposition to McMahan, that withdrawing aid and aborting a foetus are never cases of letting die. But, as we said at the outset, these are small details in McMahan’s monumental book, and they are not among the most important issues for his subject of the *ethics* of killing. For, as we have indicated, even if McMahan were right against us, that some active withdrawals of life-supporting aid qualify as instances of letting die, the moral significance claim and its euthanasia corollary would remain dubious. (It should also be remembered that McMahan does not commit himself to these doctrines; in fact, he rejects that the euthanasia corollary is a corollary of the moral significance claim for reasons into which we cannot here enter.) This is because, alongside instances of letting die that are wrong, there would still be permissible killings. But we have put these moral doctrines under greater pressure, since our analysis of the KLD-distinction enlarges the set of killings or causings of death with types of conduct many would regard as permissible. This is, we claim, because our analysis achieves an aim we share with McMahan, namely that of devising an analysis which is wholly morally neutral.