

STUDIES IN BIOETHICS

General Editor *Peter Singer*

Studies in Bioethics is a series aimed at introducing more rigorous argument into the discussion of ethical issues in medicine and the biological sciences.

Already published

Michael Lockwood (ed.) *Moral Dilemmas in Modern Medicine*

Helga Kuhse and Peter Singer *Should the Baby Live?*

The Problem of Handicapped Infants

James Rachels *The End of Life: Euthanasia and Morality*

Peter Singer and Deane Wells *The Reproduction Revolution: New Ways of Making Babies*

THE END OF LIFE

Euthanasia and Morality

JAMES RACHELS

Oxford New York

OXFORD UNIVERSITY PRESS

Oxford University Press, Walton Street, Oxford OX2 6DP

Oxford New York Toronto

Delhi Bombay Calcutta Madras Karachi

Petaling Jaya Singapore Hong Kong Tokyo

Nairobi Dar es Salaam Cape Town

Melbourne Auckland

and associated companies in

Berlin Ibadan

Oxford is a trade mark of Oxford University Press

© James Rachels 1986

First published 1986 as an Oxford University Press
paperback and simultaneously in a hardback edition

Oxford paperback reprinted 1987, 1990

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
electronic, mechanical, photocopying, recording, or otherwise, without
the prior permission of Oxford University Press

This book is sold subject to the condition that it shall not, by way
of trade or otherwise, be lent, re-sold, hired out or otherwise circulated
without the publisher's prior consent in any form of binding or cover
other than that in which it is published and without a similar condition
including this condition being imposed on the subsequent purchaser

British Library Cataloguing in Publication Data

Rachels, James

The end of life: the morality of euthanasia.

—(Studies in bioethics)

1. Euthanasia 2. Medical ethics

I. Title II. Series

174'.24 R 726

ISBN 0-19-217746-X

ISBN 0-19-286070-4 Pbk

Library of Congress Cataloging in Publication Data

Rachels, James. 1941-

The end of life.

(Studies in bioethics)

Includes bibliographical references and index.

1. Euthanasia. I. Title. II. Series.

R726.R33 1986 174'.24 B5-15361

ISBN 0-19-217746-X

ISBN 0-19-286070-4 (pbk.)

Printed in Great Britain by

Biddles Ltd

Guildford and King's Lynn

CONTENTS

Introduction	1
1. The Western tradition	
The origins of the tradition	7
The doctrine of innocence	11
The importance of being human	13
Intentional killing	15
Other views about euthanasia	17
2. The sanctity of life	
The Eastern tradition	20
Preliminary objections to the traditional views	23
A new understanding of the sanctity of life	24
The moral rule against killing	27
Some practical implications	28
Morality and religion	36
3. Death and evil	
Death and suffering	39
Hedonism	45
The concept of a life	49
4. 'Innocent humans'	
The case of Baby Jane Doe	60
The fundamental issue	62
Subnormal lives	64
'An innocent human'	67

7 ACTIVE AND PASSIVE EUTHANASIA

Killing and letting die

The idea that it is all right to allow patients to die is an old one. Four centuries before Christ Socrates said of a physician, with approval, 'bodies which disease had penetrated through and through, he would not have attempted to cure . . . he did not want to lengthen out good-for-nothing lives'. In the centuries that followed neither the Christians nor the Jews significantly altered this basic idea: both viewed allowing to die, in circumstances of hopeless suffering, as permissible. It was killing that was zealously opposed.

The morality of allowing people to die by not treating them has become more important as methods of treatment have become more sophisticated. By using such devices as respirators, heart-lung machines, and intravenous feeding, we can now keep almost anybody alive indefinitely, even after he or she has become a 'human vegetable' without thought or feeling or hope of recovery. The maintenance of life by artificial means is, in such cases, sadly pointless. Virtually everyone who has thought about the matter agrees that it is morally all right, at some point, to cease treatment and allow such people to die. In our own time, no less a figure than the Pope has reaffirmed the permission: Pius XII emphasized in 1958 that we may 'allow the patient who is virtually already dead to pass away in peace'. The American Medical Association policy statements quoted in the preceding chapter are in this tradition: they condemn mercy-killing, but say it is permissible to 'cease or omit treatment to let a terminally ill patient die'.

Thus the medical community embraces, as part of its fundamental code, a distinction between active euthanasia and what we might call 'passive euthanasia'. By 'active euthanasia' we mean taking some positive action designed to kill the patient; for

example, giving a lethal injection of potassium chloride. 'Passive euthanasia', on the other hand, means simply refraining from doing anything to keep the patient alive. In passive euthanasia we withhold medication or other life-sustaining therapy, or we refuse to perform surgery, and so on, and let the patient die 'naturally' of whatever ills already afflict him. It is the difference between *killing people*, on the one hand, and merely *letting people die* on the other.

Many writers prefer to use the term 'euthanasia' only in connection with active euthanasia. They use other words to refer to what I am calling 'passive euthanasia'—for example, instead of 'passive euthanasia' they may speak of 'death with dignity'. One reason for this choice of terms is the emotional impact of the words: it *sounds* so much better to defend 'death with dignity' than to advocate 'euthanasia' of any sort. And of course if one believes that there is a great moral difference between the two, one will prefer a terminology that puts as much psychological distance as possible between them. But nothing of substance depends on which label is used. I will stay with the terms 'active euthanasia' and 'passive euthanasia' because they are the most convenient; but other terms could be substituted without affecting my argument.

The belief that there is an important moral difference between active and passive euthanasia has obvious consequences for medical practice. It makes a difference to what doctors are willing to do. Consider this case: a patient dying from incurable cancer of the throat is in terrible pain that we can no longer satisfactorily alleviate. He is certain to die within a few days, but he decides that he does not want to go on living for those days since the pain is unbearable. So he asks the doctor to end his life now, and his family joins in the request.

One way the doctor might comply with this request is simply by killing the patient with a lethal injection. Most doctors would not do that, for all the reasons we have been considering. Yet, even so, the physician may sympathize with the dying patient's request and feel that it is reasonable for him to prefer death now rather than after a few more days of suffering. The active/passive doctrine tells the doctor what to do: it says that although he may not administer the lethal injection—that would be active euthanasia, which is forbidden—he *may* withhold treatment and let the patient die sooner than he otherwise would. It is no wonder that this simple

idea is so widely accepted, for it seems to give the doctor a way out of his dilemma without having to kill the patient, and without having to prolong the patient's agony.

I will argue, against the prevailing view, that active and passive euthanasia are morally equivalent—there is no moral difference between them. By this I mean that there is no reason to prefer one over the other as a matter of principle; the fact that one case of euthanasia is active, while another is passive, is not *itself* a reason to think one morally better than the other. My argument will not depend on assuming that either practice is acceptable or unacceptable. Here I will only argue that the two forms of euthanasia are morally equivalent: either both are acceptable or both are not. They stand or fall together. Of course, if you already think that passive euthanasia is all right, then you may conclude from this that active euthanasia must be all right, too. On the other hand, if you believe that active euthanasia is immoral, you may want to conclude that passive euthanasia is also immoral. Obviously, I prefer the former alternative; however, nothing in the argument of this chapter will depend on that.

Practical consequences of the traditional view

I will discuss the theoretical shortcomings of the traditional view at some length. However, I also want to emphasize the practical side of the issue. Employing the traditional distinction has serious adverse consequences for patients. Consider again the man with terminal cancer. Basically, the doctors have three options. First, they can end his life now by a lethal injection. Second, they can withhold treatment and allow him to die sooner than he otherwise would—this will take some time, however, so let us say that he would die in one day. And third, they could continue treatment and prolong his life as long as possible—say, for five days. (The exact numbers do not matter; they are merely for the purpose of illustration.) The traditional view says that the second, but not the first, option may be chosen.

As a practical matter, what is wrong with this? Remember that the justification for allowing the patient to die, rather than prolonging his life for a few more hopeless days, is that he is in horrible pain. One problem is that, if we simply withhold treatment, it will take him *longer* to die, and so he will suffer *more*, than if

we administered the lethal injection. Why, if we have already decided to shorten his life because of the pain, should we prefer the option that involves more suffering? This seems, on the face of it, contrary to the humanitarian impulse that prompts the decision not to prolong his life in the first place. I think I can understand why some people oppose euthanasia in any form—the view that prefers option three is mistaken, in my opinion, but it has a certain kind of integrity. A preference for the first option is also understandable. But the view which makes option two the top choice is a 'moderate' position that incorporates the worst, and not the best, of both extremes.

The cruelty lurking in the distinction between killing and letting die may also be illustrated by a very different kind of case. Down's syndrome (mongolism) is sometimes complicated by duodenal atresia (blocked intestine), and the unfortunate infant cannot obtain nourishment. In such cases, the parents and doctors have sometimes decided not to perform the surgery necessary to remove the blockage, and let the baby die. Here is one doctor's account of what happens then:

When surgery is denied [the doctor] must try to keep the infant from suffering while natural forces sap the baby's life away. As a surgeon whose natural inclination is to use the scalpel to fight off death, standing by and watching a salvageable baby die is the most emotionally exhausting experience I know. It is easy at a conference, in a theoretical discussion, to decide that such infants should be allowed to die. It is altogether different to stand by in the nursery and watch as dehydration and infection wither a tiny being over hours and days. This is a terrible ordeal for me and the hospital staff—much more so than for the parents who never set foot in the nursery.

This is not the account of a doctor who opposes the practice he is describing. On the contrary, Dr Anthony Shaw, the author of this account and one of the most frequently cited writers on the subject, supports the morality of letting these infants die. He is troubled only by the 'ordeal' he seems to think is necessary. But why is the ordeal necessary? Why must the hospital staff 'stand by in the nursery and watch as dehydration and infection wither a tiny being over hours and days'? What is gained from this, when an injection would end its life at once? No matter what you think of the lives of such infants, there seems to be no satisfactory answer. If you think that the babies' lives are precious and should be protected, then of

course you will oppose killing them *or* letting them die. On the other hand, if you think death is a permissible choice here, why shouldn't you think the injection at least as good as letting the infant 'wither'?

Let me mention another, even more bizarre, practical consequence of the traditional doctrine. Duodenal atresia is not part of Down's syndrome; it is only a condition that sometimes *accompanies* it. When duodenal atresia is present, a decision might be made to let the baby die. But when there is no intestinal blockage (or other similar defect requiring surgery), other Down's babies live on. Let us focus on this fact: *some Down's infants, with duodenal atresia, die, while other Down's infants, without duodenal atresia, live.* This, I wish to suggest, is irrational.

To bring out the irrationality of this situation, we may first ask *why* the babies with blocked intestines are allowed to die. Clearly, it is not because they have blocked intestines. The parents do not despair, and opt for death, over this condition which often could easily be corrected. The reason surgery is not performed is, obviously, that the child is mongoloid and the parents and doctors judge that because of *that* it is better for the child not to survive. But notice that the other babies, without duodenal atresia, are *also* mongoloid—they have the very same condition which dooms the ones with the blocked intestines—and yet they live on.

This is absurd, no matter what view one takes of the lives and potentials of such infants. Again, if you think that the life of such an infant is worth preserving, then what does it matter if it needs a simple operation? Or, if you think Down's syndrome so terrible that such babies may be allowed to die, then what does it matter if some babies' intestinal tracts are *not* blocked? In either case, the matter of life and death is being decided on irrelevant grounds. It is the Down's syndrome, and not the intestines, that is the issue. The issue should be decided, if at all, on *that* basis, and not be allowed to depend on the essentially irrelevant question of whether the intestinal tract is blocked.

What makes this situation possible, of course, is the idea that there is a big moral difference between letting die and killing: when there is an intestinal obstruction we can 'let the baby die', but when there is no such defect there is no choice to be made, for we must not 'kill' it. The fact that this idea leads to such results as deciding life or

death on irrelevant grounds is one reason, among others, why it should be rejected.

The Bare Difference Argument

The Equivalence Thesis, as I will call it, says that there is no morally important difference between killing and letting die; if one is permissible (or objectionable), then so is the other, and to the same degree. More precisely, it is a claim about what does, or does not, count as a morally good reason in support of a value judgement: the bare fact that one act is an act of killing, while another act is an act of 'merely' letting someone die, is not a morally good reason in support of the judgement that the former is worse than the latter.

It is compatible with the Equivalence Thesis that there may be *other* differences between such acts that are morally significant. For example, the family of an irreversibly comatose hospital patient may want their loved one to be allowed to die, but not killed. In that case, we have at least one reason to let the patient die rather than to kill him—the reason is that the family prefers it that way. This does not mean, however, that the distinction between killing and letting die *itself* is important. What is important is respecting the family's wishes. (It is often right to respect people's wishes even when we think those wishes are based on false beliefs.) In another sort of case, a patient with a painful terminal illness may want to be killed rather than allowed to die because a slow, lingering death would be agonizing. Here we have reason to kill and not let die, but once again the reason is not that one course is intrinsically preferable to the other. The reason is, rather, that the latter course would lead to more suffering.

I will argue that the Equivalence Thesis is true. It should be clear, however, that I will *not* be arguing that every act of letting die is equally as bad as every act of killing. There are lots of reasons, such as those I have just mentioned, why a particular act of killing may be morally inferior to a particular act of letting die, or vice versa. All I will argue is that, whatever reasons there may be for judging one act worse than another, the bare fact that one is killing, while the other is letting die, is not among them.

The Equivalence Thesis is one of those airy, abstract sorts of philosophical claims that may seem impossible to 'prove' one way

or the other. But I think it is possible to give some fairly convincing reasons for accepting it. The practical considerations adduced in the previous section should go some way towards making the thesis plausible; yet those considerations do not add up to a rigorous argument. What follows is an attempt to provide something more compelling.

In the sciences we often want to know what influence is exerted by one element of a complex situation. The familiar procedure is to isolate the element of interest by studying cases in which everything else is held constant, while that one element is varied. Children are taught this idea in school by having them perform simple experiments. For example, does the colour of a combustible material affect whether it will burn? Children can see that it does not by trying—and succeeding—to burn bits of paper of different colours. Does the presence of air affect combustion? Most of us will remember placing a candle in a bell-jar and watching it go out after the oxygen is consumed, while a similar candle outside the jar continues to burn. By varying one element, we see what difference it makes.

We may try a similar 'experiment' with the distinction between killing and letting die. We may consider two cases which are exactly alike except that one involves killing where the other involves letting die. Then we can ask whether this difference makes any difference to our moral assessments. It is important that the cases be *exactly* alike except for this one difference, because otherwise we cannot be confident that it is *this* difference which accounts for any variation in the assessments. Consider, then, this pair of cases:

Smith stands to gain a large inheritance if anything should happen to his six-year-old cousin. One evening while the child is taking his bath, Smith sneaks into the bathroom and drowns the child, and then arranges things so that it will look like an accident. No one is the wiser, and Smith gets his inheritance.

Jones also stands to gain if anything should happen to his six-year-old cousin. Like Smith, Jones sneaks in planning to drown the child in his bath. However, just as he enters the bathroom Jones sees the child slip, hit his head, and fall face-down in the water. Jones is delighted; he stands by, ready to push the child's head back under if necessary, but it is not necessary. With only a little thrashing about, the child drowns all by himself, 'accidentally', as Jones watches and does nothing. No one is the wiser, and Jones gets his inheritance.

Now Smith killed the child, while Jones 'merely' let the child die. That is the only difference between them. Did either man behave better, from a moral point of view? Is there a moral difference between them? *If the difference between killing and letting die were itself a morally important matter, then we should say that Jones's behaviour was less reprehensible than Smith's.* But do we want to say that? I think not, for several reasons.

First, both men acted from the same motive—personal gain—and both had exactly the same end in view when they acted. We may infer from Smith's conduct that he is a bad man, although we may withdraw or modify that judgement if we learn certain other facts about him, for example, that he is mentally deranged. But would we not also infer the very same thing about Jones from his conduct? And would not the same further considerations also be relevant to any modification of that judgement?

Second, the *results* of their conduct were the same—in both cases, the cousin ended up dead and the villain ended up with the money.

Third, suppose Jones pleaded, in his defence, 'After all, I didn't kill the child. I only stood there and let him die.' Again, if letting die were in itself less bad than killing, this defence should have at least some weight. But—morally, at least—it does not. Such a 'defence' can only be regarded as a grotesque perversion of moral reasoning.

Thus, it seems that when we are careful not to smuggle in any further differences which prejudice the issue, the bare difference between killing and letting die does not itself make any difference to the morality of actions concerning life and death. I will call this the 'Bare Difference Argument'.

Now it may be pointed out, quite properly, that the cases of euthanasia with which doctors are concerned are not like this at all. They do not involve personal gain or the destruction of normal, healthy children. Doctors are concerned only with cases in which the patient's life is of no further use to him, or in which the patient's life has become a positive burden. However, the point will be the same even in those cases: the difference between killing and letting die does not, *in itself*, make a difference, from the point of view of morality. If a doctor lets a patient die, for humane reasons, he is in the same moral position as if he had given the patient a lethal injection for humane reasons. If the decision was wrong—if, for example, the patient's illness was in fact curable—then the decision

would be equally regrettable no matter which method was used to carry it out. And if the doctor's decision was the right one, then the method he used is not itself important.

Counter-arguments

Our argument has brought us to this point: we cannot draw any moral distinction between active and passive euthanasia on the grounds that one involves killing while the other only involves letting someone die, because that is a difference that does not make a difference, from a moral point of view. Some people will find this hard to accept. One reason, I think, is that they fail to distinguish the question of whether killing is, in itself, worse than letting die, from the very different question of whether most actual cases of killing are more reprehensible than most actual cases of letting die. Most actual cases of killing are clearly terrible—think of the murders reported in the newspapers—and we hear of such cases almost every day. On the other hand, we hardly ever hear of a case of letting die, except for the actions of doctors who are motivated by humanitarian concerns. So we learn to think of killing in a much worse light than letting die; and we conclude, invalidly, that there must be something about killing which makes it *in itself* worse than letting die. But this does not follow, for it is not the bare difference between killing and letting die that makes the difference in these cases. Rather, it is the other factors—the murderer's motive of personal gain, for example, contrasted with the doctor's humanitarian motivation, or the fact that the murderer kills a healthy person while the doctor lets die a terminal patient racked with disease—that account for our different reactions to the different cases.

There are, however, some substantial arguments that may be advanced to oppose this conclusion. Here are three of them:

1. The first counter-argument focuses specifically on the concept of *being the cause of someone's death*. If we kill someone, then we are the cause of his death. But if we merely let someone die, we are not the cause; rather, he dies of whatever condition he already has. The doctor who gives the cancer patient a lethal injection will have caused his patient's death, whereas if he merely ceases treatment, the cancer and not the doctor is the cause of death. According to some thinkers, this is supposed to make a moral difference.

Ramsey, for example, urges us to remember that 'In omission no human agent causes the patient's death, directly or indirectly.' And, writing in the *Villanova Law Review*, Dr J. Russell Elkinton said that what makes the active/passive distinction important is that in passive euthanasia 'The patient does not die from the act [that is, the act of turning off a respirator] but from the underlying disease or injury.'

This argument will not do, for two reasons. First, just as there is a distinction to be drawn between being and not being the cause of someone's death, there is also a distinction to be drawn between letting someone die and not letting anyone die. It is certainly desirable, in general, not to be the cause of anyone's death; but it is also desirable, in general, not to let anyone die when we can save them. (Doctors act on this precept every day.) Therefore, we cannot draw any special conclusion about the relative desirability of passive euthanasia just on these grounds.

Second, the reason we think it is bad to be the cause of someone's death is that we think death is a great evil—and so it is. However, if we have decided that euthanasia, even passive euthanasia, is desirable in a given case, then we have decided that in *this* instance death is no greater an evil than the patient's continued existence. And if this is true, then the usual reason for not wanting to be the cause of someone's death simply does not apply. To put the point just a bit differently: There is nothing wrong with being the cause of someone's death if his death is, all things considered, a good thing. And if his death is *not* a good thing, then *no* form of euthanasia, active or passive, is justified. So once again we see that the two kinds of euthanasia stand or fall together.

2. The second counter-argument appeals to a favourite idea of philosophers, namely that our duty not to harm people is generally more stringent than our duty to help them. The law affirms this when it forbids us to kill people, or steal their goods, but does not require us in general to save people's lives or give them charity. And this is said to be not merely a point about the law, but about morality as well. We do not have a strict moral duty to help some poor man in Ethiopia—although it might be kind and generous of us if we did—but we *do* have a strict moral duty to refrain from doing anything to harm him. Killing someone is a violation of our duty not to do harm, whereas letting someone die is merely a failure

to give help. Therefore, the former is a more serious breach of morality than the latter; and so, contrary to what was said above, there is a morally significant difference between killing and letting die.

This argument has a certain superficial plausibility, but it cannot be used to show that there is a morally important difference between active and passive euthanasia. For one thing, it only seems that our duty to help people is less stringent than our duty not to harm them when we concentrate on certain sorts of cases: cases in which the people we could help are very far away, and are strangers to us; or cases in which it would be very difficult for us to help them, or in which helping would require a substantial sacrifice on our part. Many people feel that, in *these* types of cases, it may be kind and generous of us to give help, but we are not morally required to do so. Thus it is felt that when we give money for famine relief we are being especially big-hearted, and we deserve special praise—even if it would be immodest of us to seek such praise—because we are doing more than we are, strictly speaking, required to do.

However, if we think of cases in which it would be very easy for us to help someone who is close at hand and in which no great personal sacrifice is required, things look very different. Think again of the child drowning in the bathtub: *of course* anyone standing next to the tub would have a strict moral duty to help the child. Here the alleged asymmetry between the duty to help and the duty not to do harm vanishes. Since most of the cases of euthanasia with which we are concerned are of this latter type—the patient is close at hand, it is well within the professional skills of the physician to keep him alive, and so on—the alleged asymmetry has little relevance.

It should also be remembered, in considering this argument, that the duty of doctors towards their patients *is* precisely to help them; that is what doctors are supposed to do. Therefore, even if there were a general asymmetry between the duty to help and the duty not to harm—which I deny, and which I will discuss in more detail in the next chapter—it would not apply in the special case of the relation between doctors and their patients.

Finally, it is not clear that killing such a patient *is* harming him, even though in other cases it certainly is a great harm to someone to kill him. For we are going under the assumption that the patient

would be no worse off dead than he is now (otherwise, even *passive* euthanasia would be unthinkable); and if this is so, then killing him is not harming him. For the same reason we should not classify letting such a patient die as a failure to help. Therefore, even if we grant that our duty to help people is less stringent than our duty not to harm them, nothing follows about our duties with respect to killing and letting die in the special case of euthanasia.

3. The third counter-argument appeals to a consideration that has often been mentioned by doctors. Allowing a patient to die is, normally, a rather impersonal thing, in the sense that the physician does not feel 'involved' in the death—the cancer, or whatever, causes the death, and the doctor has nothing to do with it. So, there is no reason for him to feel guilty or responsible for the death. But if the physician were to give a lethal injection, *he* would be responsible, and feelings of guilt would be inevitable.

I do not wish to minimize the importance of the psychological situation in which doctors and other health-care professionals may find themselves. No doubt, many people who are comfortable enough letting die would find it psychologically impossible to kill—they just couldn't bring themselves to do it, and if they did, they would be haunted by feelings of remorse. But, important as this is for the people involved, we should be careful not to infer too much from it. We are trying to figure out whether mercy-killing is *wrong*, and whether it is *morally different* from letting die. So, we should ask: If someone feels guilty about mercy-killing, is that evidence that it is wrong? Or, if someone feels guiltier about mercy-killing than about letting die, is that evidence that it is worse?

Guilt feelings may, of course, be irrational. Someone may *feel* guilty even when he has not done anything wrong. Thus, we should not conclude that something is bad simply because someone feels, or would feel, guilty about it. We must *first* decide whether the conduct is wrong, on the basis of objective reasons; and then, if it is wrong, we may view the feelings of guilt as justified. But if it is not wrong, the feelings of guilt are irrational and we may encourage the person suffering them not to feel so bad. At any rate, feelings of guilt and the judgement of real guilt are different matters, and we cannot validly argue that a form of conduct *is* wrong, or that one type of behaviour *is* worse than another, because of feelings of guilt or innocence. That gets things the wrong way round.

The physician's commitments

Some people find it especially difficult to accept the idea of *physicians* engaging in active euthanasia. Doctors, they remind us, are dedicated to protecting and preserving life; that is their special task. Thus we should not expect *them* to kill, regardless of whatever might be right for the rest of us. Passive euthanasia, however, is another matter; since it only involves withholding pointless treatment, there is nothing in the physician's special position to rule it out.

This idea has been used by some philosophers as the basis for a qualified defence of the active/passive distinction. It is said that the distinction is important for doctors, because of their special role, regardless of whether it is important for anyone else. We need to ask, then, whether there is anything in the doctor's position that makes it impossible for him to accept active euthanasia. According to this argument, the doctor has some sort of special commitment, which the rest of us do not have, which makes the ethics of his position different.

Is this true? Everything turns on the nature of the physician's commitment. Exactly what kind of commitment is it? It might be a moral commitment—a matter of what physicians believe to be morally right—or it might be some sort of professional commitment, having to do with their role in society. Professional commitments and moral commitments are very different, and so we should consider them separately. Therefore, let's look at them one at a time.

Moral commitments

Many doctors certainly do believe that active euthanasia is immoral. Indeed, that is one reason the medical community so firmly rejects the practice. However, we should remember that anyone, including doctors, might have moral beliefs that are mistaken. To discover the truth, we must look at the arguments that can be given for and against active euthanasia; if better reasons can be given for it than against it, then it is morally acceptable, regardless of what doctors (or anyone else) might think—and of course if better reasons can be given against it, it is wrong regardless of what anyone thinks. But the fact that someone believes something is wrong never entails, by itself, that it *is* wrong. And so the fact that doctors *believe* active euthanasia to be wrong cannot, by

itself, justify the conclusion that it is wrong for them to practise it.

When thinking about this point, it is easy to fall into a certain confusion. Suppose someone mistakenly believes that something is wrong—he believes it is wrong, when in fact it is perfectly all right. If he goes ahead and does that thing, even though he believes it is wrong, he is certainly open to criticism—you may think him in some sense a morally defective person, for he should not have done what he thought was wrong. This is where the confusion can slip in. When we say 'He should not have done it' we do not mean that what he did was wrong. In *fact*, what he did was perfectly all right. Despite this, he behaved badly because he did what he *thought* was wrong. Thus, if doctors believe active euthanasia is wrong, we can say that, in this sense, they shouldn't practise it. But this will not mean that, if they did practise it, they would be doing the wrong thing.

What doctors believe is also relevant in another way. It is in general true that, other things being equal, people should be allowed to follow their own consciences. We should not, without very strong reasons, compel people to do what they think is wrong, *even if their beliefs are mistaken*. Otherwise, we do not respect their autonomy as rational beings. Thus, if a doctor believes that something is wrong, he should be permitted to refrain from it. For example, many doctors believe that abortion is immoral, and so they should not be (and in fact they are not) required to perform abortions, even though the procedure is legal and accepted by other doctors without qualm. The same might be true of active euthanasia: those physicians who disapprove of it should not have to engage in it. But it does not follow that other doctors, who take a different view, should be forbidden, and so it does not follow that it would be wrong for the medical profession in general.

Therefore, if we focus on the question of *moral* commitment, the argument we are considering fails. There is nothing in the idea of a doctor's moral commitments to support the notion that doctors are precluded from accepting active euthanasia—unless, that is, active euthanasia is objectively wrong, in which case *everyone's* moral commitments ought to forbid it.

Professional commitments

There are two ways in which doctors might be professionally committed against active euthanasia. First, it might be that doctors

pledge themselves to shun it, by subscribing to an explicit professional code of conduct. For a long time the Hippocratic Oath was taken to be such a code, although now it seems to have become more a historical relic than an actual guide. (The oath forbids abortion, for example.) More recently, the American Medical Association's policy statements have condemned active euthanasia—but it is not the purpose of those statements to bind physicians, and as we have noted, various of their provisions are regularly ignored in hospitals. Therefore, if the 'professional commitment' against mercy-killing is supposed to be in virtue of a pledge to an explicit code, there doesn't seem to be such a code. (And even if there were, advocates of active euthanasia could argue that the code should be changed.)

There is, however, another possibility. Perhaps physicians are committed to certain things, not by having taken a specific pledge, but *simply in virtue of being physicians*. Roger Rigterink, a philosopher who defends the general argument we are considering, suggests this when he says, '*The point of medicine is to preserve human life whenever it occurs,*' and '*A profession can hardly authorize an activity that is antithetical to its basic function*' (italics added). The idea is that there is something in the very conception of what it is to be a physician that rules out killing patients.

To evaluate this suggestion, let us consider a parallel argument drawn from another area of life. Suppose someone argued that, while it can sometimes be right to destroy an automobile, it can never be right *for a mechanic* to do such a thing. After all, the whole point of automobile mechanics is to repair cars and make them serviceable. In destroying a car, a mechanic would be going against the very nature of his profession. So, if an automobile is beyond salvaging, it may be acceptable for the owner to junk it, but he cannot expect the mechanic to have any part of such a thing.

Obviously, this is a silly argument. But why? It isn't because there is something less noble about automobile mechanics than there is about doctors; nor is it silly because cars lack special moral worth. (There is no 'sanctity of automotive life'.) It is a bad argument because the concept of a profession cannot be used to show that it is wrong for a professional to do something that falls outside that concept. Mechanics fix cars when they can be fixed; if, in

consigning a jalopy to the junk-heap, he isn't acting 'as a mechanic', what of it? Similarly, if a doctor, in practising active euthanasia, wasn't acting 'as a doctor', what of it?

It might be objected that, in engaging in mercy-killing, the physician isn't merely doing something *outside* his professional role; he is doing something *incompatible* with it. However, the same can be said about the auto mechanic. If the point of that profession is 'to repair cars and make them serviceable'—and isn't that its point?—then it is equally incompatible with auto mechanics to junk cars. Nevertheless, we would think it very strange for a mechanic to insist that he can do nothing to help us junk cars without violating his calling.

Suppose, however, we admit for purposes of argument that the nature of medicine does somehow imply that those engaged in it cannot be involved in mercy-killing. Would that mean we must meekly accept the implied conclusion? No, for consider this: we can define a *different* profession, very much like medicine, but called (perhaps) 'smedicine'. Smedicine, as we will define it, is the profession which does everything it can to treat illness, cure disease, and repair the human body, so long as there is any point to it; but, when the possibility of a meaningful life is gone, smedicine helps to make the passage to death as easy as possible. We could argue that medicine, which (we are assuming) precludes this latter kind of help, is morally defective, and should be abandoned, to be replaced by the better practice of smedicine. If I thought that the concept of medicine precluded mercy-killing, that is exactly what I would argue, for 'medicine', thus conceived, would be forbidding its practitioners from doing what is in many instances the morally right thing.

There is, therefore, nothing in the physician's commitments that leads to the conclusion that the active/passive distinction is somehow valid 'for him'. If it is in general an unsound distinction, then it is as much unsound for him as for anyone.

Thomson's objection

The Bare Difference Argument relies on a certain method of reasoning, and some philosophers have suggested that this method is not sound. Judith Jarvis Thomson has urged that *something* must be wrong with this way of reasoning, because it leads to patently

absurd conclusions. To demonstrate this, she offers an argument that is parallel to the one involving Smith and Jones, but which is obviously unsound:

Alfrieda knows that if she cuts off Alfred's head he will die, and wanting him to die, cuts it off; Bertha knows that if she punches Bert in the nose he will die—Bert is in peculiar physical condition—and, wanting him to die, punches him in the nose. But what Bertha does is surely every bit as bad as what Alfrieda does. So cutting off a man's head isn't worse than punching a man in the nose.

She concludes that, since this absurd argument doesn't prove anything, the Smith-and-Jones argument doesn't prove anything either.

If Thomson were right, we would have to scuttle the Bare Difference Argument and look elsewhere for support for the Equivalence Thesis. But I don't think she is right: the Alfrieda-and-Bertha argument is not absurd, as strange as it is. A little analysis shows that it is a sound argument and that its conclusion is true. The analysis is a bit tedious, but it is worth doing, for it clarifies the nature of the Bare Difference Argument and confirms its soundness.

We need first to notice that the reason it is wrong to chop someone's head off is, obviously, that this causes death. (I am setting aside secondary reactions having to do with messiness.) The act is objectionable because of its consequences. Thus, a different act with the same consequences may be equally objectionable. In Thomson's example, punching Bert in the nose has the same consequences as chopping off Alfred's head; and, indeed, the two actions are equally bad.

Now the Alfrieda-and-Bertha argument presupposes a distinction between the act of chopping off someone's head, and the results of that act, the victim's death. (It is stipulated that, except for the fact that Alfrieda chops off someone's head, while Bertha punches someone in the nose, the two acts are 'in all other respects alike'. The 'other respects' include the act's consequence, the victim's death.) This is not a distinction we would normally think to make, since we cannot in fact cut off someone's head without killing him. Yet in thought the distinction can be drawn. The question raised in the argument, then, is whether, *considered apart*

from their consequences, head-chopping is worse than nose-punching. And the answer to *this* strange question is No, just as the argument says it should be.

The conclusion of the argument should be construed like this: the bare fact that one act is an act of head-chopping, while another act is an act of nose-punching, is not a reason for judging the former to be worse than the latter. At the same time—and this is perfectly compatible with the argument—the fact that one act causes death, while another does not, *is* a reason for judging the former to be worse. Thomson has specified, however, that in the cases of Alfrieda and Bertha there is no difference in this regard either; and so their acts turn out to be morally equivalent. So be it.

The parallel construal of the conclusion to the Smith-and-Jones argument is: the bare fact that one act is an act of killing, while another act is an act of letting die, is not a reason for judging the former to be worse than the latter. At the same time—and this is perfectly compatible with that argument—the fact that an act (of killing, for example) prevents suffering, while another act (of letting die, for example) does not, is a reason for preferring the former. So once we see exactly how the Alfrieda-and-Bertha argument *is* parallel to the Smith-and-Jones argument, we find that Thomson's argument is, surprisingly, quite all right. Therefore, it provides no reason for doubting the soundness of the style of reasoning employed in the Bare Difference Argument.

The Compromise View

Some philosophers concede that, in the case of Smith and Jones, there is no moral difference between killing and letting die; but they continue to maintain that in the euthanasia cases the distinction is morally important. Thus, it is suggested that the Bare Difference Argument commits an elementary fallacy—the fallacy of hasty generalization. It leaps from one example, in which the distinction appears to be unimportant, to the general conclusion that the distinction is *never* important. But why should this be so? Why should the only options be that the distinction is *always* important, or *never* important?

Perhaps the truth is simply that the difference between killing and letting die is sometimes morally important, and sometimes not, depending on the particular case you choose to think about.

I will call this the Compromise View. It is the most appealing alternative to the view I am defending, and of all possible views of the matter it most closely conforms to our pre-reflective intuitions. The Compromise View allows us to look at cases one at a time, and decide in each case whether the difference between killing and letting die is significant. What could be more reasonable?

In fact, I will argue, such a procedure is not reasonable at all. Logic requires that the distinction be always, or never, important. There is no middle ground. This may sound unattractively dogmatic. Nevertheless, it follows from some inescapable principles of reasoning.

The crucial question is this: Is it possible for a fact sometimes to count as a good reason in support of a moral judgement, and sometimes not? Imagine what it would be like if this were possible. I tell you that John is a bad man because he is stingy and a liar; you then observe that Frank is also a stingy liar, and so you conclude that he is a bad man as well. But I object, and say that although stinginess and dishonesty count against John, the same does not apply to Frank. Frank, I say, is a splendid fellow. I am *not* saying that, despite being a stingy liar, Frank has other qualities that compensate for this. I am saying something more radical—I take the fact that John is stingy and dishonest to be a good reason for judging him to be a bad man; but I do not take the fact that Frank has these qualities to be a reason for judging him badly. I am holding these qualities against John, but I am not holding these very same qualities against Frank at all, not even as something that needs to be compensated for.

Or, to take a different example: I say that you ought to vote for Brown, a candidate for public office, because she favours gun control. You point out that Black, her opponent, also favours gun control; therefore you say that I have not given any reason for preferring Brown over Black. But again, I object, and say that Brown's position is a reason in her favour, but Black's identical position is not a reason in Black's favour.

Surely, in both these cases, I am inconsistent. It would be perfectly all right to argue that Brown, but not Black, should be elected for other reasons. Perhaps we know more about them than that they both favour gun control: perhaps we also know that both candidates are strong supporters of affirmative action

programmes, and that Brown has greater experience in dealing with governmental matters. Then we might tabulate what we know like this:

Reasons for voting for Brown

Brown favours gun control.

Brown supports affirmative action programmes.

Brown has greater experience.

Reasons for voting for Black

Black favours gun control.

Black supports affirmative action programmes.

We certainly may conclude that, all things considered, one ought to vote for Brown. In every respect save one, they are equally good candidates; and in the one respect in which they differ, Brown has the edge. So she is the candidate of choice, at least on this information. What we *cannot* do, without violating the requirement of consistency, is say that Brown's position on gun control (or her position on affirmative action) goes on the list in her favour, but that Black's position doesn't even go on the list.

There is a formal principle of reasoning involved here; 'formal' because it is a principle of logic that everyone must accept regardless of the content of his or her particular moral code. Let *A* and *B* stand for any actions, and let *P* stand for any property of actions. Then:

Principle I

If the fact that *A* has *P* is a morally good reason in support of the judgement that *A* ought (or ought not) to be done, and *B* also has *P*, then that is also a reason, of equal weight, for the judgement that *B* ought (or ought not) to be done.

The act 'voting for Brown' has the property 'being a vote' for someone who favours gun control', and the act 'voting for Black' has the same property; so, if this provides a reason for voting for Brown, it also provides a reason for voting for Black.

The following corollary covers cases in which the merits of two acts are being compared:

Principle II

If the fact that *A* has *P* and *B* has *Q* is a morally good reason for preferring *A* over *B*, then if *C* has *P* and *D* has *Q*, that is a reason of equal weight for preferring *C* over *D*.

Thus, if we say that Brown's experience as compared with Black's is a reason for preferring Brown, then in any other similar election between an experienced and an inexperienced candidate, we must say this is a reason for preferring the former.

Superficially, there appear to be some counter-examples to this principle—that is, examples which show the principle to be unsound. Suppose we prefer Brown's experience over Black's in one election, but in another election we think it important to bring in a candidate from outside government—'fresh blood', as it is called. Doesn't this mean that, in one instance, we are taking experience as desirable, and in another instance taking it as undesirable—and can't this be all right? Of course this is all right—it certainly isn't irrational—but it does not violate the principle. It only appears that we have violated the principle because we have not specified the reasons accurately. The reason we sometimes prefer an 'inexperienced' candidate is not that he is inexperienced. It is because we think it likely that he will have a fresher approach, be more open to new ideas and be less bound by the mistakes of the past. And these are *always* good qualities in a candidate. Other apparent counter-examples to the principle may be explained away in a similar manner.

Now, with these principles in mind, let us return to the distinction between killing and letting die. You will recall that the Equivalence Thesis is a thesis about what does, or does not, count as a morally good reason in support of a value judgement. It says: the fact that one act is an act of killing, while another act is an act of letting someone die, is not a morally good reason in support of the judgement that either act is preferable to the other. The Compromise View, on the other hand, says that in some cases the distinction may be important, while in other cases it is not. In other words, the Compromise View implies this:

In some cases, the fact that *A* is letting die, while *B* is killing, is a morally good reason for preferring *A* over *B*. But in other cases, the fact that *C* is letting die, while *D* is killing, may *not* be a morally good reason for preferring *C* over *D*.

And this violates Principle II. This is what I meant when I said that the Compromise View is inconsistent with sound principles of moral reasoning. The reasonable-sounding compromise offered by

this view is not tenable; and so we are stuck with the radical-sounding alternatives: either the distinction between killing and letting die is always important, or it never is.

What do advocates of the Compromise View actually say? Philippa Foot proposes a version of the Compromise View. She asks, 'When is this distinction morally relevant?' and answers 'in cases in which rights are in question'. It may violate a person's rights to kill him, she says, even when in the same circumstances it would not violate his rights to let him die. She adds that 'permission may make all the difference in such a case . . . if someone gives us the right to kill him, for his own good, by seriously consenting to the action, it then makes no moral difference whether we do kill him or rather allow him to die.'

Foot thinks it is 'obvious' that there is sometimes a moral difference between killing and letting die, especially 'when one thinks about the crucial question of rights'. She suggests that those who defend the Equivalence Thesis are confused about what they are denying; they 'seem simply to have misunderstood the position of their opponents'.

It would not be surprising if there were some misunderstanding here. Often, when people take sides on complex issues, and are determined to defend their views, misunderstanding occurs. People become more interested in scoring debating points, and proving themselves right, than in patiently analysing issues. Rather than being immune from this tendency, philosophers often seem to be among the worst offenders.

I can think of one way to bring the two sides on this issue a little closer together. I (and other defenders of the Equivalence Thesis) can easily concede that, in some cases, it may be permissible to let die but not to kill. Likewise, in other cases it may be permissible to kill but not to let die. This is perfectly compatible with the Equivalence Thesis. All the Equivalence Thesis requires is that, in such a case, it is some *other* feature of the case that makes the difference.

Permission may, indeed, be a crucial matter, as Foot says. Suppose I give you permission to let me die, but say that I do not want to be killed. Then it might be all right for you to let me die but not to kill me. (If it is the other way round—I give permission to be killed, but don't want to be allowed to die—the reverse may be true.) However, the *reason* one is permissible, but not the other, will

not be that one is killing and the other is letting die. The reason will be, simply, that I permitted one but not the other.

Perhaps the Equivalence Thesis is misunderstood by those who reject it. Foot emphasizes that 'If one may, in particular circumstances, allow a man to die it does not *follow* that one may also kill him, even for his own good.' She apparently thinks this is a telling point against those of us on the other side of the issue. But it isn't; this is something we all agree on. Suppose, again, that someone wants to be allowed to die but does not want to be killed. In those 'particular circumstances', one may allow him to die; but, as she says, 'it does not *follow* that one may also kill him'. The Equivalence Thesis does not imply otherwise. The Equivalence Thesis only says that the reason one course, but not the other, is permissible, isn't simply the intrinsic 'moral importance' of the difference between killing and letting die. In this case that difference is *correlated* with another difference (between permission and objection) and this other difference is, indeed, morally important.

I do not know whether clearing away such misunderstandings, and exposing areas of common ground, will lead to agreement. It should at least reduce the extent of the disagreement, and clarify the differences that remain. Complete agreement is probably too much to hope for; that is rare in any branch of philosophy, and even rarer when philosophical theses lie so close to moral practice.

8 FURTHER REFLECTIONS ON KILLING AND LETTING DIE

The status of intuitions

In the previous chapter I argued that there is no moral difference between killing and letting die. In addition to its obvious importance for medical ethics, this thesis has other implications: it implies, for example, that our duty towards starving people is greater than we might have realized. In this chapter, I take up this issue, and present additional arguments for the Equivalence Thesis. This will also be a convenient place to comment on some questions of method in ethical theory.

Although we do not know exactly how many people die each year of malnutrition and related health problems, the number is very high, in the millions. The most common pattern among children in poor countries is death from dehydration caused by diarrhoea brought on by malnutrition. In 1983, not a particularly bad year, James Grant, Executive Director of the United Nations Children's Fund (UNICEF), estimated that 15,000 children were dying in this way every day. That comes to 5,475,000 children annually. Even if Grant's estimate was high by a factor of three, it is still a staggering number of deaths—and it includes only one way of death, among only one class of victims.

By giving money to support famine-relief efforts, each of us could save at least some of these people. By not giving, we let them die. The Equivalence Thesis suggests a harsh judgement about this behaviour, for it says there is no moral difference between letting die and killing. The question of famine relief provides another important test of this idea, for even if one is convinced that there is no difference in medical contexts—that active and passive euthanasia are morally the same—one might still balk at the idea in connection with famine relief. Moreover, there is likely to be more resistance to this application of the Equivalence Thesis because it