End-Of-Life Decisions and Moral Psychology: Killing, Letting Die, Intention and Foresight

Word Counts:

Abstract: 219

Main text: 6866

Endnotes: 215

Discipline: Philosophy

Topic Area: End of Life Decisions

Keywords: Intention; Double Effect Principle; Euthanasia; Withholding Treatment;

Terminal Care

Charles Douglas, B Med, FRACS

Senior Lecturer, Clinical Ethics and Health Law, The University of Newcastle,
Callaghan, NSW, 2308, AUSTRALIA and

PhD Candidate and Associate, Centre for Value, Ethics and the Law in Medicine, The
University of Sydney, NSW, 2006, AUSTRALIA

Phone 0400 390 748

Abstract

In contemplating any life and death moral dilemma, one is often struck by the possible importance of two distinctions; the distinction between killing and “letting die”, and the distinction between an intentional killing and an action aimed at some other outcome that causes death as a foreseen but unintended “side-effect”. Many feel intuitively that these distinctions are morally significant, but attempts to explain why this might be so have been unconvincing. In this paper, I explore the problem from an explicitly consequentialist point of view. I first review and endorse the arguments that the distinctions cannot be drawn with perfect clarity, and that they do not have the kind of fundamental significance required to defend an absolute prohibition on killing. I go on to argue that the distinctions are nonetheless important. A complete consequentialist account of morality must include a consideration of our need and ability to construct and follow rules; our instincts about these rules; and the consequences (to the agent and to others) that might follow if the agent breaks a good general rule, particularly if this involves acting contrary to moral instinct. With this perspective, I suggest that the distinctions between killing and letting die and between intending and foreseeing do have moral relevance, especially for those involved in the care of the sick and dying.

Introduction

Central to the debate about euthanasia, at least in a medical context, are two controversial distinctions. The first is the distinction between killing and letting die, which I will call the DKL. Turning off a ventilator (with the inevitable death of a ventilator-dependent patient) is, for example, contrasted with the administration of a lethal injection to the same patient in the same circumstances. The second distinction is that between death intended as a means to some other end, and death caused as a foreseen but unintended side effect of an action performed with a different aim. A doctor who administers a lethal injection with the intention of killing her terminally ill patient, as a means to relieving suffering, is contrasted with a doctor who administers large doses of analgesics and sedatives, with the sole aim of relieving suffering but in the knowledge that death may be hastened. This latter distinction is at the heart of the so-called “Doctrine of Double Effect” but I will be concerned in this article with the central distinction between intending and foreseeing (DIF) rather than a full explication of the doctrine.

There are those who believe that there are *no* exceptions to the moral rule against killing the innocent (even in the case of euthanasia where the “victim” clearly wants death to come quickly). I will refer to this view as “absolutist”. The DIF and DKL are clearly critical to absolutism. It is more or less universally accepted that in some circumstances we must let patients die; and that in some circumstances, we must accept that a patient’s death may be hastened by justifiable action, particularly the administration of sedatives and analgesics.ⁱ If it is *sometimes* morally acceptable to let a patient die, but *never* to kill him, there must be a recognisable (and morally relevant) difference between killing and letting die. If it is sometimes morally

acceptable to cause the death of a patient as a side effect of one's actions, but not to kill him, then there must be a recognisable (and morally relevant) difference between intending death and "merely" foreseeing death as a result of one's actions.

The corollary is that there are arguments against absolutism, beginning with a critique of the DKL and/or the DIF, that segue into arguments in favour of euthanasia (Kuhse 1987; Rachels 1994; Tooley 1994). The purpose of this article is not to challenge or endorse conclusions about euthanasia, but to reassess the role of these two distinctions. I firstly review what I believe are valid arguments about the problems with the DKL and DIF, specifically that they cannot be clearly described, and that they are not morally relevant "in themselves". I endorse the view that they cannot sustain an absolutist morality; but I hope to show that the DIF and DKL are nonetheless morally relevant distinctions.

A caveat is necessary. Rights of patients are central to end-of-life decisions, at least insofar as doctors generally have a duty to provide care to patients who want to be treated, an obligation to withhold treatment from patients who refuse it, and an obligation to consider what a non-competent patient might or might not have consented to, had she been competent. I set aside these issues in this paper, not because they are unimportant, but in order to focus on the issue of intentional killing once consent issues have been addressed to the extent that they can.

Distinctions that are not morally relevant ““in themselves””

The distinction between killing and letting die

There is, clearly enough, a commonsense view that some outcomes are made to happen (by some kind of *positive* involvement of an agent) and some outcomes are allowed to happen (by some kind of *negative* involvement of the agent). The DKL is a special case of this general distinction. However it has not proven possible, to date, to establish a functional definition of positive and negative agency in a way that corresponds with the commonsense view consistently (Howard-Snyder 2002). A particularly significant difficulty pertains to the removal of a metaphorical safety-net, an example of which is described in the book *Touching the Void* (Simpson 1989). Mountaineer Simon Yates deliberately cut the rope supporting his companion Joe Simpson when he could no longer support them both. Had Simpson died, it is not clear whether we should say that he had been killed or merely allowed to fall to his death. The same problem applies to the withdrawal of mechanical respiratory support in the ventilator-dependent patient. Indeed, withdrawal of ventilation might look like a killing *and/or* a letting die, depending on the circumstances (McMahon 1994; Thomson 1999).

Obviously this lack of clarity alone is a major problem for the absolutist point of view. But there is a further difficulty for the absolutist, and that is to prove the moral significance of the positive-negative distinction. While it is easy enough to think of paired examples that suggest a moral difference (Trammell 1994), it is open to argument in these cases that the apparent moral difference is due to factors usually but not always associated with positive-negative agency rather than due to the positive-negative distinction itself (Bennett 1998, 75-77). On the other hand, it is possible to

describe positive-negative pairs in which there seems to be *no* moral difference. James Rachels has given what is probably one of the most famous examples in the medical literature, that of the equally evil Smith and Jones, two men who both plan to drown their young cousins in the bath for personal gain (Rachels 1994). Smith succeeds; Jones doesn't need to, because (just before he can act) his cousin accidentally slips and drowns while Jones watches. Moral intuition supports the notion that there is nothing to prefer in Jones' negative agency, and that the bare difference between killing and letting die is unimportant. Merits of Rachels' example include that the events are at least plausible in real life, and that the behaviours of Smith and Jones would count as killing and letting die respectively according to almost any description of the distinction. Whether paired examples of this sort can ever be conclusive is arguable (Bennett 1998, 77), but it seems fair to say that the burden of proof — to describe what the doing-allowing difference really is, and why it matters — remains with the absolutist, and is yet to be met (Howard-Snyder 2002).

The distinction between intending and foreseeing

The DIF rests on the idea that we can usually recognise that our actions have multiple possible outcomes, some of which are desirable and others about which we may feel neutral, or may even prefer not to cause at all, but accept as necessary if we are to achieve our primary aims. If a surgeon performs an operation, her aim is to heal (and presumably also to earn a living), but an anticipated side-effect is that she will cause pain. The pain is foreseen (perhaps even with certainty) but is not in common sense aimed at or intended. The distinction thus stated looks reasonably sound, and at a glance it seems to have obvious moral significance – foreseen harm must be acceptable, whereas intended harm (a surgeon trying to cause pain) must not.

Problems arise when we consider an interim “bad” outcome (M) that we can use as a *means* to a “good” final outcome (E) – if this is the case, we may have to aim at M in order to bring about M (and thus E), and in common sense we *do* then intend M even though we don’t desire it (except for its instrumental value). For example, a woman may be genetically predisposed to breast cancer. With her consent, a surgeon may remove the woman’s healthy breasts, as a means to preventing her premature death. The immediate outcome of the surgery — being without breasts — is in commonsense bad, but this is precisely the outcome that the surgeon and patient must aim for. As Elizabeth Anscombe says of “justified amputations”, “there really is here a case in which *an* intrinsic evil is chosen for the sake of a more valuable good” (Anscombe 2001). The moral significance of the DIF in this context is underwhelming. The fact that the surgeon *intends* something bad for its instrumental value surely doesn’t make her action morally bad. Nor is it obvious that this form of surgery is morally inferior to, say, cosmetic breast surgery, in which the bad outcomes (such as pain) are merely foreseen.

However, we should consider the possibility that the intention-foresight distinction is pivotal only in the area of killing, as Anscombe has suggested (Anscombe 2001, 64). The following three examples are fairly typical of the way in which the DIF is supposed to have moral significance, where death is an outcome. I have described the examples in terms of what is permissible, in accordance with an absolutist view. In the first pair of examples, (call them Wartime 1 and 2) it could be morally acceptable for “our side” to drop bombs on military targets, repeatedly, knowing (but not intending) that many civilians will be killed (W1); but it could never be morally

acceptable for us to deliberately target a civilian population (as perhaps happened in Hiroshima and Nagasaki) as a *means* to forcing the enemy into an early surrender (W2), even if this may save thousands of lives compared to W1 over the course of the war. In a second pair of scenarios (Foetus 1 and 2), it would be permissible for a doctor to perform a hysterectomy on a pregnant woman with cancer of the uterus, the intention being to remove the cancer, with the death of the foetus inside merely foreseen (F1). However, it would be impermissible to abort the pregnancy of a woman with severe heart disease (to protect her from the dangers of pregnancy, and thus save her life), as the intention would apparently be to kill the foetus as a *means* to saving the mother (F2). In the third pair of cases (Palliation/Euthanasia 1 and 2), a doctor may give very large doses of sedatives to a patient with terminal cancer and uncontrollable agitation, foreseeing that the patient's death is *certain* to be hastened (PE1)ⁱⁱ; but a soldier whose companion is pinned under the wreckage of a burning vehicle (dying an agonising death) may not shoot him in the head (PE2) in order to put him out of his misery.

Now, many people will probably find the first case in each of these pairs more morally palatable than the second. Whether they would accept a prohibition on the second cases is another matter. But I have so far ignored a further difficulty, that in hard cases such as these the ascription of intention is contentious. In particular, there is this problem: we can *define* intention in a very fine-grained way, so that each outcome is made up of a number of others, and any undesirable death can be seen to be a side-effect of something else. This is most easily illustrated (I think) in the foetal example above. In F2, the doctor can perform an abortion by merely *removing* the foetus from the mother, foreseeing its death occurring in that process. The death of

the foetus in itself is not required as a means to saving the mother, and if the foetus were old enough to survive its removal from the womb, the doctor would not make a further attempt to kill it. This form of reasoning can be extended to any hard case. The soldier in PE2 does not need to kill his companion to end his suffering, only to make him immediately unconscious (with death being a foreseeable but more or less inevitable consequence of the bullet fired into his head). And, perhaps most bizarrely, as Jonathon Bennett has provocatively suggested (Bennett 1981, 67), the dropping of a bomb on civilians might only be done to cause the civilian bodies to disintegrate temporarily, to make them look dead, with permanent cessation of life being something that followed as an inevitable and almost immediate side-effect.

Of course the notion of making people look dead by dismembering them is obviously far too close to killing for most of us (including Bennett) to take seriously the idea that a so-called “terror bomber” does not intend to kill – this defies common sense. But there is nothing clear and distinct that allows us to say that an intention to dismember or to separate a foetus from a womb, or to fire a bullet into a head *is* an intention to kill – there is only a rough notion of “felt closeness” first described by Phillipa Foot (Foot 2001). And “closeness” has proven to be as good as it gets when it comes to judging, in a way that accords with commonsense, whether a killing is intentional or not (Bennett 1998; Foot 2001; Delaney 2008).

Now, if separating a foetus from its mother’s womb (to unburden the mother) is close enough to a killing, then the act of Antarctic explorer Captain Oates, in walking out of his tent into a blizzard (to unburden his companions) should surely also count as an intended death. As O’Keeffe says (O’Keeffe 1984, 357), “If he did not knowingly

and intentionally go to his death, he would not have accomplished his purpose”. But here we have a case that seems morally justifiable and perhaps even commendable. The bare fact that Oates may have *intended* to kill himself (as a means to an end) seems to carry no weight at all in this case. Of course, there are also those prepared to deny that Oates intended to kill himself (Holland 1975), but I think this is a case of “fine-graining” that goes against the commonsense view. We cannot just decide about intention on a case by case basis because there is obviously a temptation for those who have embraced an absolutist point of view to adjust their intuitions about intention and their intuitions about which deaths are morally acceptable to ensure a perfect match (McIntyre 2006; Delaney 2008).

In summary then, the DIF is problematic in sustaining an absolute prohibition on killing. A fine-grained description of the distinction offends common sense by putting too much on the foreseen (and thus permissible) side; while a common-sense view of the distinction is at best too vague to offer guidance and at worst prohibits morally acceptable actions.

Throwing out the cousin with the bathwater

If the DKL and DIF cannot be drawn in such a way that they can be shown to have moral relevance *in themselves*, do these distinctions have any role in moral thinking? Can they now be dismissed? Some philosophers seem to have implied so, at least within the medical context. James Rachels used the example described above of Smith and Jones to raise questions about the use of the DKL in a medical context. Rachels wrote that “killing is not in itself any worse than letting die; if my contention

is right, it follows that active euthanasia is not any worse than passive euthanasia” (Rachels 1994). Now, on a narrow reading of Rachels, he is just saying that active euthanasia *in itself* is not any worse than passive euthanasia. If that is what he means, the logic is obviously correct. But it is hard to see that this claim is particularly helpful as a moral statement. To say that an action of kind A is not *in itself* worse than an action of kind B is indeed a denial of a strict deontological view of morality; but it doesn’t prove that we should abandon rule-following behaviour all together.

To illustrate this point further, imagine a variation in Rachels’ thought experiment. Just as Jones notices his cousin slip in the bath, and resolves to let him drown, the child’s mother enters the bathroom from another doorway, sees the child underwater and picks him up. Then another accident occurs: the mother slips, hitting the child’s head on the bathtub. The child suffers a fatal injury, and there is nothing Jones can do to prevent this (even if he wanted to). So the outcome is the same as in Smith’s case; and Jones’ and Smith’s intentions and actions are the same as in the original example. I think that Rachels would have to say that there is still no morally relevant difference between the scenarios. But Jones has performed a “nothing” (his agency being irrelevant to the death) whereas Smith has killed his cousin. It is patently absurd (or at least pointless) to conclude from this that killing is no worse than “not killing”. Such a contrived example surely has no wider significance at all.

What is concerning, then, is that Rachels’ argument is misleading because it encourages a broader reading, something like this: “killing is not in itself any worse than letting die...it follows that you should think of active euthanasia as being no worse than passive euthanasia”. Such an argument is logically wrong. Moral

psychology, our thinking about what we should do in a situation that we perceive as moral, is complex: it involves amongst other things, pattern recognition, rule following, habitual behaviour and reacting to instinct or intuition. It does not follow from Rachels' example that I should cultivate a belief that killing and letting die are in some broader moral sense equivalent; that as soon as I have resolved not to interfere in a person's dying that I may as well kill her, all else being equal.

Why the DKL and DIF are morally relevant distinctions

I have argued that just because the DKL and DIF have not been shown to be morally relevant distinctions *in themselves*, we cannot conclude that they are, in a broader sense, morally irrelevant. I now wish to go further and make the positive argument that they are indeed morally relevant distinctions, at least within the medical context (although the argument is generalisable). In doing so, I rely on the assertion that the DKL and DIF are recognisable distinctions. Notwithstanding the fact that intentions may be multiple and to a certain extent ambiguous (Quill 1993; Douglas, Kerridge and Ankeny 2008), and that many doctors have an exaggerated belief about the respiratory depressant effects of narcotics (Sykes and Thorns 2003), doctors do generally know when the dose of a drug is *obviously* in excess of that required to relieve symptoms, and is being used with an intention to kill (Douglas, Kerridge et al. 2008, 393). And doctors can recognise the difference between introducing an intervention that kills, and withdrawing or withholding an intervention or treatment that may be sustaining life. The argument is not diminished if there are occasional counter-examples.

There are three reasons why the DKL and the DIF are morally relevant distinctions. All three reasons relate to the importance of a general (as opposed to an absolute) rule against killing, and the roles of the DKL and DIF in sustaining such a general rule. The first reason relates to the *probability* of immediate death following an action of a given kind. The second reason relates to the moral psychology of rule following and the way that we think about exceptions to a rule. The third reason relates to instincts about the DKL and the DIF.

1. The probability of immediate death

The likelihood that death will result from an action is uncontroversially a relevant factor in moral evaluation of that action (Tooley 1994, 107; Bennett 1998, 224). In thought experiments, probabilities are (sometimes bizarrely) imagined to be equal in contrasting doing/allowing and intending/foreseeing pairs, so that philosophers can try to focus their attention on the *intrinsic* importance (or lack thereof) of the DKL and the DIF. However, it is fairly obvious that an attempt to kill will *generally* be associated with an increased probability of death compared to “refraining from saving someone’s life” (Tooley 1994, 107). And in real life, the agent who aims to kill is more likely to cause death than the agent who does not so aim (Foot 2001, 147).

The outcomes of most courses of action are by no means certain, and as real life moral agents, we evaluate actions (in our deliberations) according to the *estimated* likelihood that they will achieve a given outcome. Therefore the *kinds* of actions that we should contrast when we evaluate the doing/allowing distinction are not killing and letting die (descriptions which presuppose the outcome as certain), but as Tooley has implied, “attempted killing” and “refraining from saving life”.ⁱⁱⁱ When evaluating

the intending/foreseeing distinction we should be contrasting attempted killing with actions that are not aimed at killing (but carry some risk of death).

To take this argument into the medical context, what should interest us is not just a special case (letting die) but *all* cases of limiting (withdrawing or withholding) treatment. Without begging the question about its moral relevance, we cannot presuppose that a doctor who limits treatment does even *expect* to cause a patient's death immediately (let alone intend or desire to do so). Now admittedly there are cases where a doctor must surely expect that a patient will die promptly, as in the withdrawal of ventilatory support from a patient who cannot breathe independently. But if we are only interested in comparing killing with limiting treatment in settings where both actions will result in more or less *certain* and *immediate* death, then our conclusions will also be limited to those settings. "Letting die", defined according to the apparent belief that death will certainly follow, is a course of action open to doctors in very few circumstances indeed. "Limiting treatment" is an action open to doctors in every setting, *as is (attempted) killing*, and these are the two kinds of actions that we should be examining for their moral relevance. We can still use the acronym DKL (distinction between killing and limiting treatment).

Similarly, "causing death as a *reasonably certain* foreseen but unintended side effect of treatment" is an almost vanishingly rare event in medicine, even in a palliative care setting (Sykes and Thorns 2003). If we are to consider the wider moral relevance of the DIF, the *kind* of action that we should be examining (on the other side of the line from killing) is medical treatment in which death is foreseen as a *possible* outcome

(with probability being anything from 0 to 100%). I will call this “giving treatment with a known risk of death” for want of a less clumsy expression.

To see the significance of this, imagine the following common clinical scenarios.

Where it is decided not to treat a patient’s cancer with adjuvant chemotherapy (thus dropping the chance of cure from 93% to 90%, but avoiding the morbidity of chemotherapy), there is nothing consequentially close to a killing occurring. Where a cure is not considered possible, but “palliative” chemotherapy is contemplated, a decision to forgo such chemotherapy may result in a loss of some time, but the patient may still have months to live. Where a potentially curative operation is performed, with a known mortality rate of 5%, a foreseeable death *may* occur, but 100 of such operations will lead to very different consequences indeed compared to 100 consecutive (attempted) killings. Where a decision is made to forgo antibiotics, a patient may die within days, but may also recover without antibiotics – this may *possibly* turn out to be a letting die, but there is a nonzero probability that the patient will survive for months or *years*, and the action therefore differs substantially from a killing. Where a terminally ill patient with poorly controlled pain is given an infusion of morphine in hospice, the morphine is likely to have a negligible influence on the timing of death (Fohr 1998). Where a *very* agitated patient with a terminal illness requires very high doses of sedatives to control his agitation, inducing unconsciousness, and intravenous fluids are concomitantly withdrawn, the patient is highly likely to die soon; as is the case when an obviously ventilator-dependent patient is extubated. Now, only in these last two cases do the probabilities suggest that there is something similar to a killing going on in simplistic outcome terms. Perhaps (setting aside my second and third arguments that follow), in these last two cases and

none of the others, one could invoke the “moral irrelevance” of the DIF or the DKL and say that the patients could as well be killed – but since we have already accepted that death will come quickly and certainly in these latter cases, killing is neither worse *nor better* (in terms of direct consequences).

Of course, the situations in which voluntary euthanasia is advocated are precisely those in which killing and limiting treatment are *not* consequentially similar, it being argued that killing may be *better* than limiting treatment, as there may be less time for a patient to suffer (Rachels 1994). Imagine for example a patient who has terminal cancer, and an expected survival of six months, who requests euthanasia, and refuses any form of medical treatment. Respecting the wish for euthanasia leads to an immediate death; respecting the refusal of treatment results in a death over six months (or maybe a week if the patient refuses to eat or drink). The euthanasia advocate would ask “why is the high likelihood and immediacy of death a bad thing?”. That is not what is at issue here: the point is that it is obviously a *morally relevant* thing. (Attempted) killing is problematic first and foremost *because* it leads to an immediate and certain loss of life in (almost) all circumstances in which it occurs. And the greater the difference in consequential terms of a killing from a limiting of treatment — or a treatment with known risk of death, in probability-outcome terms — the greater the burden on the agent in making the decision.

It may be argued that I have only pointed out a *general* difference in consequences between killing, on the one hand, and limiting treatment and giving treatment with a known risk of death on the other, and that this tells us nothing about individual cases. This is true. But most consequentialists will accept that we need morally relevant

guides and that these guides require that we distinguish some *kinds* of action from others (Mackie 1977, 155; Hare 1993). Doctors make many life and death decisions, and they must have some strategy for making these decisions. At one extreme there would be *just* this rule: do whatever provides the best chance of prolonging life. There is clearly a consensus that this is not a good moral rule to follow in all instances. But to what strategy do we fall back, when this first strategy no longer seems appropriate? Are there any rules that a doctor can follow, once it is no longer obvious that more life is a good thing? I would suggest that there are three possible rules as follows:

- Rule 1: Limit treatment – don't kill, but withdraw or withhold treatments that seem to be causing more harm than good.
- Rule 2: Treat symptoms – don't kill, but if the symptoms justify it, allow treatments where the risks of death are significant.
- Rule 3: Decide on a case by case basis whether the patient should be killed or not.

The third rule leaves the agent with an enormous burden of judgment. Even if the main criterion for determining whether or not to kill someone were voluntariness, and even if this could be determined in a reliable and sensitive way, we would still have a problem. The number of patients who ask to be killed is a tiny fraction of the number of patients whose lives are nearing an end (van der Maas et al. 1991), and in all other cases, doctors are faced with decisions about what to do at the end of life. As Hare says, doctors need general principles based on the ordinary run of cases, not on the hard cases or philosophers' examples (Hare 1993). Between curing and killing there are recognisable *kinds* of action described by Rules 1 and 2 above. It seems perverse

to suggest that these kinds of action are not different to killing in a morally relevant way.

This is not to say that there are not enormous challenges within the context of limiting treatment or treating symptoms, or that these decisions are not contentious. Much of the criticism of the DKL has focused on the “selective” non-treatment of neonates with, for example, spina bifida or Down’s syndrome, and an inter-current but fairly easily treatable condition (Kuhse 1987; Rachels 1994, 113). It would seem that some of these are indeed cases where non-treatment is almost certain to be followed fairly quickly by death *and* where death is in commonsense the intention of the limitation of treatment. As such, these cases more closely resemble killings, and not surprisingly, *this* kind of letting die is highly contentious (Steinbock 1994). There will always be hard cases. But the vast majority of “lettings die” in medical practice (in my own experience) are nothing like this – they are daily events that have more in common with the Yates-Simpson case, where we are cutting a metaphorical rope when the burden of holding on is no longer sustainable. And sometimes, as Yates did, we get a pleasant surprise.

2. The problem of making exceptions to a rule

I have suggested that it is appropriate to make use of Rules 1 and 2 above as *general* rules. But just as the alcoholic cannot abstain without cultivating a belief that every single drink is harmful (even though this is not true *in itself*), we cannot follow a rule without cultivating *some* sense that individual exceptions to that rule are “wrong”.

The first argument then, was that killing is *generally* morally different to limiting treatment or treating with a known risk of death because it *generally* leads to a more

certain and immediate death. The second argument is that an *individual* instance of killing is morally different to, and indeed more problematic than, an individual instance of limiting treatment or treating with a known risk of death, even where the immediate consequences are the same or seem to favour the killing, simply because the killing is recognisably a killing, and because we know that the performance of this *kind* of action has the potential to undermine the *general* rule against killing.

3. Morally significant instincts

The fact that we have certain moral intuitions does not mean that these intuitions reflect objective moral truths (Mackie 1977, 38). But this point needs to be distinguished from a separate one about intuition. To the extent that an intuition is widely shared and deeply ingrained (call it an instinct), there is an overwhelming argument that it *must* be accommodated in our morality (see, for example, Bennett 1981, 89). The point can be made more obvious by considering our most basic instincts. Consider the following proposition:

Humans have an instinct for self-preservation

P1

P1 is a claim for which there would surely be little dispute, at least as a generalisation. It is a fact with a plausible biological explanation, but with no obvious moral justification. And yet, we don't doubt that the difference between living and dying is a morally relevant difference, even if it is "just" a biological instinct that we prefer to live. The instinct is a fundamental building block for morality; it is what leads us to count death, generally, as a "bad" outcome, and a life saved as a "good" outcome when we consider the consequences of our actions.

Now, consider 2 further propositions:

Humans have an instinct that killing is worse than letting die P2

Humans have an instinct that intentionally causing death is worse than P3
causing death as an unintended side-effect of one's actions

There is empirical evidence to support both propositions; evidence to suggest that the intuitive appeal of the DKL and the DIF is widely shared, and not a matter of culture or upbringing (Hauser 2006). And there are very plausible reasons why we should have these instincts — an outcome is vastly more likely to occur if one tries to bring it about than if one does not. So trying to cause harm is worse than not trying, and acting to bring about harm is worse than not acting. These are beliefs that we might reasonably have inherited, not just learned, and they may therefore be beliefs that are “hard-wired” and relatively rigid, not “soft-wired” and plastic.

Of course P2 and P3 are not synonymous with normative moral judgements (for example, that killing *is* worse than letting die). If P2 and P3 are true, and we do have an instinctual belief in the moral significance of the DKL and DIF, it does not follow that the instincts so described should, or inevitably will, overwhelm our moral reasoning. But if they are hard-wired, we are likely to *feel* that the DKL and DIF must be “true”. Furthermore, these instincts will influence the way that we feel about any life and death decision that we must make, *over and above* the influence of simple consequential considerations (as consequentialism is traditionally conceived). Our instincts may then be a very important part of our rule-following behaviour.

Firstly, there may be distress to an agent who decides to break the rule, and to those who observe and know about the case, simply because reasoning alone cannot easily make the killing feel right. Secondly, to the extent that one can resolve any psychological distress associated with the killing, the initial importance of the rule may be very significantly undermined. In other words, the instinctual element of the rule against killing might reinforce obedience to the rule *and* magnify the undermining effects of a breach of the rule. As a colleague remarked of her reluctance to be involved in an ostensibly justifiable case of feticide, “I’m not sure what it is — perhaps a fear of releasing the evil within”.

Discussion

The DKL, the DIF and the euthanasia debate

It should be clear enough from my discussion that I do not believe an absolute deontological view of morality is tenable, and therefore, the moral acceptability of a particular case of euthanasia must always be open to question. I cannot, for example, see a principled argument to leave a person to burn to death when trapped in a fire if the chances of rescue are negligible (as in PE2 above). The traditional consequentialist arguments in such a case are overwhelmingly in support of shooting the trapped person; the rule-breaking consequences of doing so are minimal because of the isolated nature of the case; and we are only left with the moral instinct against killing, which I would argue that we *should* overcome in such a difficult circumstance. While one can indeed argue that shooting the trapped person is not an action intended to kill (but only to make unconscious), such a claim in *this* case would I think be disingenuous.

It is also not my purpose here to make an argument with respect to the legalisation of medically assisted death. I would only make the following comments: the fact that there are problems with the DKL and the DIF, most obvious when examining hard cases, does count as an argument against an absolute moral prohibition on euthanasia. On the other hand, the fact that the DKL and DIF are morally relevant distinctions in anything but their bare bones does count as an argument in favour of maintaining a general moral prohibition on euthanasia in a medical context. How we enforce the general rule while dealing sympathetically with its justifiable exceptions is, I think, the obvious dilemma that we face.

I have discussed the importance of a moral instinct against killing, but such an instinct does not necessarily trump the interests of an individual patient, or the doctor's other intuitions. There may be competing moral instincts, such as an instinct to relieve suffering, and where suffering seems to resist all available medical means, a doctor may feel compelled *toward* euthanasia. Instincts may conflict, as do reasons. But a complete consequentialist account of euthanasia should at least *include* a consideration of the concerns that Gillett has raised (Gillett 1994, 325) regarding the instincts of doctors and what it means to them to act as (compassionate) killers.

Moral instinct, “closeness” and the tragedy of hard choices

The arguments that I have raised about moral instinct have general philosophical implications. Whether an instinctual element to beliefs about the DKL and DIF does make them actually “intrinsically” morally relevant I think is a matter of definition. But it is interesting to speculate on the role that instinct might play in the convictions

of moral philosophers. Strong intuitions in favour of these distinctions is widespread (Howard-Snyder 2002; McIntyre 2006), and moral philosophers with quite different overall perspectives have argued that deontological constraints based more or less roughly on the DKL and the DIF must be right (for example, Quinn 1994; Quinn 2001). Yet to date none have been able to advance a convincing theoretical explanation (Scanlon 2008, 2). Thomas Nagel flirted with the notion of biological instinct as the basis for the appeal of the DIF: “it is tempting to think that the whole thing is a kind of moral illusion resulting (from) innate psychological dispositions”, before apparently rejecting that idea because the intuition seems to present itself “as the apprehension of *truth*, not just as a psychological inhibition” (Nagel 2001). But it is hard to see how one can make such a judgment, anymore than I can separate my *genuine* love of life from a “mere” survival instinct.

I have pointed out above that an absolutist view about the wrongness of intentional killing cannot be sustained without invoking “closeness” between the desired outcome and the death to which that outcome is linked. Closeness seems a fairly fuzzy concept to reason with, but from the point of view of moral psychology and moral instinct, closeness makes some sense. A death may *feel* worse (and therefore may have more significant consequences to the agent, both immediately and in terms of future behaviour) to the extent that it is *closely* tied to the agent’s intentional action and is therefore *recognisably* a killing kind of action in a primitive sense. There is empirical evidence to support this (Greene 2005).

Now the traditional consequentialist response to such an argument may be this: “such feelings are not morally relevant — get over your instincts, and do what is right”. But

I think that this is short sighted. These feelings *matter*, because they have psychological consequences for the agent *and* because the act of quelling them has consequences for future behaviour. This is well illustrated I think by the example of Sophie's choice cited by Magnusson (Magnusson 2006), but varied slightly. Sophie was the fictional character invited by the Nazi's to choose just one of her children or lose both to the gas chambers. The fact that she *chose* is a matter of obvious moral relevance. We can argue about whether, in choosing to keep her son, Sophie killed her daughter; or (merely) let her die or unintentionally caused her death, but I think these are the kind of philosophical questions to which there are no logically correct answers. To put the matter beyond doubt for the strict deontologist, we can imagine that Sophie was handed a pistol and asked to shoot one child in order to save the other, an obviously "impermissible" action. If choosing one child's life over another is not too close to a killing, then firing a bullet into her head clearly is. Now what of the consequentialist point of view? As it is traditionally conceived, and perhaps caricatured, there should be no hesitation; Sophie should choose one life over two (even if that means flipping a coin before deciding whom to shoot). But a consequentialism more in keeping with commonsense recognises something peculiarly tragic about the situation simply because Sophie must become the chooser, the doer, the one who *aims at* her own child's death. It comes as no surprise that the fictional Sophie is ultimately harmed, not just because she has lost a child, but because *her* agency is involved; and that the consequences of her choice amount to more than "one life lost, one saved".

Conclusion

Between an indefensibly rigid deontological view that killing is always wrong and a naïve consequentialism in which there is no morally relevant difference between

killing and letting die, and no morally relevant difference between intending and foreseeing death, there is room for a consequentialist morality sympathetic to deontology and virtue ethics. The moral acceptability of a particular instance of killing may be open to question, *but* in any contemplated circumstance we must take into account, not only the *likelihood*, number and quality of lives lost, but also the psychological consequences to the agent of overriding a basic moral instinct, and the consequences to the agent and to the community (in terms of future behaviour) of allowing an exception to a basic moral rule. These considerations are generally quite different to those faced when a doctor decides, in a medical context, to withdraw or withhold treatment or to give sedatives or analgesics to a terminally ill patient.

Notes

ⁱ It is often implied that death may be a very likely or even inevitable consequence of the administration of analgesics and sedatives to terminally ill patients (see for example Kuhse 1987, 86). There is a widespread consensus in modern palliative care circles that this is rarely the case (Ashby 1997; Fohr 1998; Sykes and Thorns 2003). On the other hand, it is conceded that *sometimes* palliative treatment does indeed hasten death (Sykes and Thorns 2003, 343), which clearly leaves a question to be answered by the absolutist.

ⁱⁱ As pointed out in note i above, such a situation is exceedingly rare. The claim of certainty is generally unrealistic and is only made here in order to maintain some symmetry between the pairs; as is argued in the main text, differences in certainty are in reality an important part of the distinction between intending and foreseeing.

ⁱⁱⁱ In this light, a flaw in Rachels' intuition test is apparent: Jones sets out to kill his cousin. Only luck intervenes, but if it had not, Jones *would* have proceeded to try to kill his cousin, and would most likely have succeeded. Jones' ultimate action may be an "allowing", but in some sense it is an attempted killing, and it naturally affects our intuitions in a similar way to a killing.

Acknowledgements

I would like to thank Associate Professor Ian Kerridge and Associate Professor Rachel Ankeny for extensive comments on the ideas in this paper and reviews of the draft manuscript

Competing Interests

None

References

- Anscombe, E. 2001. Medalist's address: Action, intention, and 'double effect'. In *The doctrine of double effect: Philosophers debate a controversial moral principle*, ed. P. A. Woodward. Notre Dame: University of Notre Dame Press. Originally published in *Proceedings of the American Catholic Philosophical Association* 56, 1982.
- Ashby, M. 1997. The fallacies of death causation in palliative care. *Med J Aust* 166: 176.
- Bennett, J. 1981. Morality and consequences. The Tanner Lectures on Human Values, II. Salt Lake City: University of Utah Press.
- Bennett, J. 1998. *The act itself*. Oxford, Oxford University Press.
- Delaney, N. F. 2008. Two cheers for "Closeness": Terror, targeting and double effect. *Philosophical Studies* 137(3): 335-367.
- Douglas, C., I. Kerridge and R. Ankeny. 2008. Managing intentions: The end-of-life administration of analgesics and sedatives, and the possibility of slow euthanasia. *Bioethics* 22(7): 388-396.
- Fohr, S. A. 1998. The double effect of pain medication: Separating myth from reality. *Journal of Palliative Medicine* 1(4): 315-328.
- Foot, P. 2001. The problem of abortion and the doctrine of double effect. In *The doctrine of double effect: Philosophers debate a controversial moral principle*, ed. P. A. Woodward. Notre Dame: University of Notre Dame Press. Originally published in *Virtues and Vices and Other Essays in Moral Philosophy, 1978*, University of California Press.
- Gillett, G. 1994. Killing, letting die and moral perception. *Bioethics* 8(4): 312-328.
- Greene, J. 2005. Cognitive neuroscience and the structure of the moral mind. In *The innate mind*, ed. P. Carruthers, S. Laurence and S. Stich. Oxford: Oxford University Press.
- Hare, R. M. 1993. Medical ethics: Can the moral philosopher help? In *Essays on bioethics*. Oxford: Oxford University Press. Originally published in 1977. *Philosophical Medical Ethics: Its Nature and Significance*, ed. S. Spicker and H.T. Engelhardt Jr: 49-96. Reidel/Kluwer.
- Hauser, M. 2006. *Moral minds*. New York: HarperCollins Publishers.
- Holland, R. F. 1975. Suicide. In *Moral problems : A collection of philosophical essays* ed J. Rachels. New York: Harper & Row. Originally published in *Talk of God (Royal Institute of Philosophy Lectures Volume 2, 1967-8)*.
- Howard-Snyder, F. 2002. Doing vs allowing harm. In *Stanford Encyclopedia of Philosophy Fall 2008 Edition*, ed. E. N. Zalta. <http://plato.stanford.edu/entries/doing-allowing/>. Accessed 1 June 2009.
- Kuhse, H. 1987. *The sanctity-of-life doctrine in medicine: A critique*. Oxford: Clarendon Press.
- Mackie, J. L. 1977. *Ethics: Inventing right and wrong*. London: Penguin.
- Magnusson, R. S. 2006. The devil's choice: Re-thinking law, ethics, and symptom relief in palliative care. *The Journal of Law, Medicine & Ethics* 34(3): 559-569.
- McIntyre, A. 2006. Doctrine of double effect. In *The Stanford Encyclopedia of Philosophy Summer 2006 Edition*, ed. E. N. Zalta.

- <http://plato.stanford.edu/archives/sum2006/entries/double-effect>. Accessed 1 June 2009.
- McMahon, J. 1994. Killing, letting die, and withdrawing aid. In *Killing and letting die*, ed. B. Steinbock and A. Norcross: 250-279. New York: Fordham University Press. Originally published in *Ethics* 103: 250-279.
- Nagel, T. 2001. Agent-relative morality. In *The doctrine of double effect: Philosophers debate a controversial moral principle*, ed. P. A. Woodward. Notre Dame: University of Notre Dame Press. Originally published in "The Limits of Objectivity," *The Tanner Lectures on Human Values, 1980*, University of Utah Press.
- O'Keefe, T. M. 1984. Suicide and self-starvation. *Philosophy* 59(229): 349-363.
- Quill, T. E. 1993. The ambiguity of clinical intentions. *N Engl J Med* 329(14): 1039-1040.
- Quinn, W. S. 1994. Actions, intentions and consequences: The doctrine of doing and allowing. In *Killing and letting die*, ed. B. Steinbock and A. Norcross New York: Fordham University Press. Originally published in *Philosophical Review*, 98.
- Quinn, W. S. 2001. Actions, intentions, and consequences: The doctrine of double effect. In *The doctrine of double effect: Philosophers debate a controversial moral principle*, ed. P. A. Woodward. Notre Dame: University of Notre Dame Press. Originally published in *Philosophy and Public Affairs*, 18, 334-351.
- Rachels, J. 1994. Active and passive euthanasia. In *Killing and letting die*, ed. B. Steinbock and A. Norcross: 112-119. New York: Fordham University Press. Originally published in *N Engl J Med*, 292, 78-80.
- Scanlon, T. M. 2008. *Moral dimensions*. Cambridge: Belknap Press of Harvard University Press.
- Simpson, J. 1989. *Touching the void*. New York: Harper and Row.
- Steinbock, B. 1994. The intentional termination of life. In *Killing and letting die*, ed. B. Steinbock and A. Norcross. New York: Fordham University Press. Originally published in *Ethics in Science and Medicine*, 6.
- Sykes, N. and A. Thorns. 2003. Sedative use in the last week of life and the implications for end-of-life decision making. *Archives of Internal Medicine* 163(3): 341-344.
- Sykes, N. and A. Thorns. 2003. The use of opioids and sedatives at the end of life. *The Lancet Oncology* 4(5): 312-318.
- Thomson, J. J. 1999. Physician-assisted suicide: Two moral arguments. *Ethics* 109(3): 497.
- Tooley, M. 1994. An irrelevant consideration: Killing versus letting die. In *Killing and letting die*, ed. B. Steinbock and A. Norcross, 103-111. New York: Fordham University Press.
- Trammell, R. 1994. Saving life and taking life. In *Killing and letting die*, ed. B. Steinbock and A. Norcross, 290-297. New York: Fordham University Press. Originally published in *Journal of Philosophy*, LXXII, 131-137.
- van der Maas, P. J., J. J. van Delden, L. Pijnenborg and C. W. Looman. 1991. Euthanasia and other medical decisions concerning the end of life.[see comment]. *Lancet* 338(8768): 669-74.