

THE MORALITY OF INTERFERENCE

Christian Munthe
(University of Gothenburg)

This is a pre-print. The final version, which is the one to use for citations, is published in *Theoria*, 65 (1): 55–69. Available online at: <http://onlinelibrary.wiley.com/doi/10.1111/j.1755-2567.1999.tb00114.x/abstract>

1. Introduction

The distinction between doing harm and allowing harm plays an important role in non-utilitarian ethics.¹ It is used to resolve apparent conflicts of duties or rights (e.g., in cases when either of the alternatives open to an agent will result in someone's death). It is employed to underpin arguments to the effect that utilitarianism permits too much (e.g., killing one innocent person to save two others), and that it demands too much (e.g., giving up your personal projects in order to save starving people in Africa).² In political philosophy, the doing-allowing distinction has been used to explain the basic conditions of property-rights. It also surfaces in many

¹ I use "utilitarianism" to denote only versions of *act*-utilitarianism, according to which an action, *a*, is morally right iff. there is no alternative to *a* which would, if performed, have produced a higher balance of agent-neutral positive value over agent-neutral negative value than *a*. Rule-utilitarianism is disregarded since it is possible that such an ethical theory incorporates the distinction between doing and allowing in the system of rules which, due to the overall utility of it being generally accepted or complied to in society, it is our duty to follow, according to rule-utilitarianism. I also refrain from speaking broadly about consequentialism, since, as Philippa Foot has observed (1985, pp. 27-29), consequentialist theories may very well employ the distinction between doing and allowing, for example a so-called utilitarianism of rights (Cf. Nozick 1974, pp. 26-30) or consequentialist theories giving weight to agent-relative basic values in such a way that a harm of a certain magnitude that is *done* is counted as worse than an *allowed* harm of the same magnitude.

² Sometimes these explanations instead use the notions of positive and negative duties or rights. But the basic idea seems to be the same. See, e.g., Jonathan Glover's discussion of these concepts (1977, p. 97).

issues within applied ethics, such as the moral status of euthanasia. Finally, it is reflected in many people's common-sense opinions. For example, most people would spontaneously judge it morally worse to stab a beggar to death than to refuse him the 5\$ he asks for, thereby ensuring his death.

Several attempts have been made to explain this distinction, and what makes (or does not make) it morally significant.³ This paper will focus on *one* such idea only, namely the suggestion that the difference between doing and allowing is to be understood in terms of *interference* versus non-interference. According to this idea (from now on referred to as "the distinction"), when I *do* harm I interfere with a causal flow, which by itself would not have resulted in the harm in question, thereby making it to produce this harm. When I *allow* harm to happen, however, I rather abstain from interfering with a causal flow, which by itself leads to the harm.⁴

Much of the debate regarding this distinction has concerned its ability to match every-day uses of pairs of notions such as "active-passive", "commission-omission", "doing-allowing", "making-letting", "action-inaction" and so forth. However, notwithstanding the general interest of such investigations, I will focus more directly on the *moral* significance of the distinction. For even if, as Shelly Kagan and Warren Quinn have suggested,⁵ the distinction is unable to match perfectly our linguistic intuitions regarding the use of such pairs of words, it may nevertheless capture what is of *moral* importance in such uses. As I see it, then, the primary aim of the distinction is not to capture perfectly our linguistic intuitions about the difference between doing and allowing, but to find a *morally* significant division of *different ways of causing harm*.

Bennett and Kagan have both denied the distinction such significance. Kagan's reason for this is based on the role he assigns it to play in a certain kind of ethical theory he calls *moderate*. This role is to ground an absolute

³ See, for example, Bennet 1966; 1981; and 1995, chapters 4-8; Dinello 1971; Donagan 1977; Finnis 1973, p. 139; Foot 1967; and 1985; Gauthier 1986, pp. 203-205; Glover 1977, chapter 7; Greenwell 1975, pp. 21-22; Kagan 1988; and 1989, chapter 3; Quinn 1989; and Rachels 1975.

⁴ Bennet (1995, chapter 7) sees this idea as the primary rival to his own analysis of the doing-allowing distinction and attributes it to Alan Donagan (*op.cit.*). It is also the one judged to be the most promising in Shelly Kagan's (*op.cit.*) comprehensive critical discussion of the doing-allowing distinction. Moreover, it has been used by Gauthier (*op.cit.*) in his discussion of the 'Lockean proviso' for the original acquisition of property-rights. Foot has overcome her initial scepticism about the prospect of making sense of the doing-allowing distinction (1967), and has put forward the interference-idea as her new stand on the issue (1985).

⁵ Kagan, *op.cit.*, and Quinn, *op.cit.*

moral prohibition against doing harm – an essential part of any moderate theory – without going against any other essential feature of such a theory.⁶ Because of this, he sees it fit to place the burden of proof on those moderates who claim the distinction to be morally significant.⁷ His treatment ends with an invitation to supporters of the doing-allowing distinction in a moderate framework to explain...

...why doing is morally more significant than allowing, why doing harm is more significant than doing good, and why harm-doing generates agent-relative reasons rather than agent-neutral ones.⁸

Bennett suggests that those who believe the interference idea to be morally relevant might be influenced by some confusions,⁹ and then proceeds:

With that muddle set aside, I see no way of finding moral significance in [the distinction between interference and non-interference] on our present version of it. That is not to say that it has none, [...] but I do contend that those who think that it is ever morally significant owe us reason for saying so. I cannot find that any reasons have been given.¹⁰

Both Kagan's and Bennett's reasoning rely heavily on claims regarding the proper placement of the burden of proof. Supporters of the moral significance of the distinction are told that they have to put forward some argument for their view, while those who deny the idea are not requested to substantiate this denial any further. Of course, supporters of the moral significance of the distinction owe us reason for their claim, but this does not mean that *denials* of that claim escape the need for supporting arguments. Thus, as long as no reason against the moral significance of the distinction has been presented, there is no reason to deny it such relevance.¹¹

⁶ See Kagan, *op.cit.*, chapter 1, for further explanations of the terminology and Kagan's general approach.

⁷ This is made clear from Kagan's initial discussion of the moral relevance of the distinction (*op.cit.*, pp. 111-114).

⁸ Kagan, *op.cit.*, p. 127.

⁹ Bennett, *op.cit.*, p. 119.

¹⁰ *Ibid.*, pp. 119-120.

¹¹ Furthermore, I do not want to presuppose, as Kagan does, any *special* kind of non-utilitarian theory in which the claim that the distinction is morally significant is to be incorporated. It is quite sufficient that this claim, since it puts *some* moral weight on something else than the (agent-neutral) value of consequences of actions, is inconsistent

My aim is to decrease this epistemological discrimination against supporters of the moral significance of the distinction, and present an argument which supports Bennett's and Kagan's claim *without* placing the burden of proof on their (our) opponents. I will argue that, no matter how the idea of interference is varied (within a generally plausible framework for analysing statements about interference) it is *clearly* irrelevant for moral matters. My argument for this claim is that its denial would violate a very plausible extension of the principle that 'ought' implies 'can'.

This claim does not rule out that the idea of a morally significant distinction between doing harm and allowing harm cannot be specified in some other way which could escape my argument. However, the idea of interference is the latest among such suggested specifications, and also the most intuitively plausible according to several commentators (myself included). Moreover, my discussion only regards the *basic* moral significance of the distinction – i.e., the idea that the property pinpointed by this distinction is a right-/wrong-making characteristic. My claim that it is not compatible with the suggestion that the distinction may carry some moral importance of an instrumental nature. For example, utilitarians may claim that a *belief in* the (basic) moral significance of the distinction is a part of desirable trace of character, notwithstanding the fact that this belief is mistaken.¹²

I will start out in section 2 by explaining more closely the distinction, and how it may be incorporated into a basic moral principle. In section 3, I will formulate my argument against such moral principles.

2. Interference Explained

The general idea of a difference between doing harm and allowing harm is best explained with reference to a "paradigm-example":¹³

with utilitarianism. Thus, a supporter of the moral significance of the distinction does not have to answer either of the last two of Kagan's questions unless she also wants to defend a moderate theory. Regarding the first question, it would, in my view, be close to begging the question to criticise such a supporter for not answering it in a satisfactory way.

¹² In Richard Hare's terminology, even if the distinction has no place on the *critical* level of moral thinking, it may have an important role to play on the *intuitive* level. See Hare 1981.

¹³ The point of this example is *not* to formulate what Kagan has called a *contrast argument* (see Kagan 1988) against the moral significance of the distinction, but mainly to serve as an instrument for making clear what the distinction amounts to.

Case 1: A holds B's head under water, thereby drowning her.

Case 2: D has happened to fall into the water. C can save her, but refuses. D drowns.

Any explication of the doing-allowing distinction should imply that A is doing harm, while C is allowing harm to occur (assuming that death is harmful for B and D).

The *general* idea of the difference between causing harm by interference and causing harm by non-interference is the following. In the case of interference, a causal flow which would be harmless if uninterrupted is disrupted by the agent's behaviour in a way which effects a harmful causal flow. In the case of non-interference, a causal flow which would be harmful if uninterrupted, could be disrupted by the agent's behaviour in a way which would effect a less harmful causal flow, but is in fact not so disrupted.

Consider case 1, where B initially is swimming peacefully in the pool with no or very low risk of drowning. She is thereby part of a causal flow which by itself does not have her death as an upshot. Then A suddenly appears on the scene and presses B's head under water, thereby altering this casual flow in such a way that the death of B is effected. Contrast this with case 2, where the initial causal flow of which D is a part *would* by itself effect D's death. C *could*, but does in fact *not*, behave in a way that would alter this causal flow in such a way that D does not die. The role played by C's behaviour in the causal history of D's death in case 2 is thus reversed compared to the role played by A's behaviour in the causal history of B's death in case 1. It is this difference which makes it the case that A harms B by interference, while C harms D by non-interference.

This talk about causal flows which by themselves would produce this or that upshot if only they were not disrupted by some behaviour calls for a *general* framework for talk about causality which can be used to pinpoint the relevant differences between cases such as 1 and cases such as 2. The framework that has actually been used in the literature, both by those who support the moral significance of the distinction and by its critics, is the analysis of causal statements in terms of counterfactual dependence. Simply put, in this framework, statements such as "event E causes event F" can be analysed as stating the fact that F occurs after E and that F would not have occurred if E had not occurred. This is not to say that the analysis

of counterfactuals is unproblematic.¹⁴ However, this analysis allows that hazy talk about different roles of agents' behaviours in the causal histories of upshots is put into a language which is both familiar and capable of describing the mechanisms involved in complex causal schemes such as cases 1 and 2.

It is, of course, possible that some other framework for analysing talk of causality comes to be proved to be superior to the counterfactual analysis, and that, in turn, that framework will be better suited for analysing the difference between interference and non-interference. If so, the argument I will formulate below would have to be restated in terms of that analysis. However, in the following I will simply assume a counterfactual analysis of statements of causality.

The basic form of an analysis of the difference between interference and non-interference in terms of counterfactuals looks as follows. Let *X* denote a false statement about an agent causing some avoidable harm. This agent causes the harm by *interference* if, and only if, the harm would *not* have occurred, had *X* been true. For example, if *X* had been true in case 1, B would not have drowned. The harm is instead caused by *non-interference* of the agent if, and only if, the harm *would* have occurred anyway, had *X* been true. For example, if *X* had been true in case 2, D would have drowned anyway.

In order to flesh out this formula into an actual analysis, *X* needs to be specified in a plausible way. Several suggestions are possible, but the formula just demonstrated sets one limit: *X* must not be specified such that its truth would yield the *same* upshot in cases 1 and 2. The following suggestions for specifying *X* have been put forward in the debate: 1) the agent *has never existed*,¹⁵ 2) the agent *does not exist at the time* at which the harm was actually inflicted,¹⁶ 3) the agent *is not present at the occasion* on which the harm was actually inflicted,¹⁷ and 4) the agent *is in no position to make a choice* that will affect whether or not the harm occurs.¹⁸ I will return to these suggestions when formulating my argument in section

¹⁴ See, for example, Goodman 1979, especially pp. 3-27; and Jackson 1991.

¹⁵ Kagan, *op.cit.*, p. 94.

¹⁶ Kagan, *op.cit.*, p. 96.

¹⁷ Kagan, *op.cit.*, p. 99, note 11. This suggestion is identical to Gauthier's (*op.cit.*, p. 204).

¹⁸ Bennet, *op.cit.*, p. 118. Kagan also considers a fifth specification according to which *X* says that the agent does not behave in the way he actually does (Kagan, *op.cit.*, p. 97). However, this suggestion gives the wrong implication in the paradigm-example, since it implies that the truth of *X* would effect the same upshot in case 1 and 2.

3 below.

In order to incorporate the distinction into a basic moral principle it is useful to formulate it as two separate principles:

I: An action, *a*, which effects some harm is a case of causing harm by interference if, and only if, the victim of the harm would have fared better, had *X* been true.

NI: An action, *a*, which effects some harm is a case of causing harm by non-interference if, and only if, *a* does not meet the condition of **I**.

These two principles may be employed in a basic ethical theory for making a moral difference between doing harm and allowing harm in several ways. I will start with the simplest suggestion:

MI: An act, *a*, which meets the condition of **I** is, to some extent, morally worse, *ceteris paribus*, than another act, *b*, which meets the condition of **NI**.¹⁹

The expression "to some extent" in **MI** ensures that there may be different versions of **MI**, giving different moral weight to differences of the amount of harm caused by *a* and *b* respectively. For example, while one version of **MI** may say that I am not permitted to kill the one whatever the consequences, other versions may permit me to save, say, the *ten* by killing the one. The *ceteris paribus* clause ensures that an ethical theory incorporating **MI** may also incorporate *other* right- and wrong-making characteristics (such as the intentions of the agent, special obligations, acts being performed in self-defence and so on). Thus, two ethical theories, both of which incorporate **MI**, may very well judge the relative moral seriousness of an act, *a*, which meets **I**, compared to that of another act, *b*, which meets **NI**, very differently *all things considered*.

It is this capability of **MI** to be just one part of a more complex ethical theory that makes it possible for supporters of **MI** to handle the possibility that **I** and **NI** are at odds with linguistic intuitions about doing and allowing. As I said at the beginning, this does not seem to be any problem in principle for someone who wants to put forward a basic moral principle such as **MI**, as long as this principle does not imply judgements which are

¹⁹ **MI** has many similarities to the principle formulated and rejected by Bennet (*op.cit.*, p. 118).

at odds with firmly held *moral* opinions. But this is not very probable concerning **MI**, since it does not imply any *categorical* moral judgements, but only *ceteris paribus* ones. Other right- and wrong-making characteristics such as the agent's intentions, special obligations, how difficult it would have been for the agent to act differently, the seriousness of the harm caused (including side-effects) and the degree to which the harm is certified to occur by the agent's behaviour can, in a pluralistic moral theory incorporating **MI**, pull in other directions than **MI** and thus secure that the theory yields acceptable categorical moral judgements.

Although **MI** is compatible with many different non-utilitarian ethical theories there is, however, one *minimal* principle which have to be endorsed by *anyone* who endorses an ethical theory which includes *some* version of **MI**:

MMI: An act, *a*, which meets the condition of **I** is morally worse, *ceteris paribus*, than another act, *b*, which meets the conditions of **NI**, provided that the harm caused by *b* is just as morally serious as the harm caused by *a*.

MI and **MMI** say that the question of whether or not an action causes harm by interference or by non-interference is *always* morally significant in the same way regardless of the presence of other morally relevant factors. A basic moral principle employing the distinction may be more complex, however. For example, one idea is that facts about interference – non-interference become significant for the comparison of the relative moral wrongness of two actions only if some *other* difference between these action holds (or, alternatively, that its significance becomes *stronger or weaker* when such an additional difference holds). In other words, **MI** and **MMI** may be supplemented with conditions such as: "provided that *a* is performed from evil intentions and *b* is not" or "provided that *a* and *b* are not cases of self-defence".²⁰ However, since these kind of complication do not affect the power of my argument against **MMI** and **MI** (soon to be set forth), I will proceed with the simple versions of these principles and only make a brief comment below on how my argument can be extended to rule out also more complex versions of **MI** and **MMI**.

²⁰ C.f. Kagan 1988. I am grateful to Gustaf Arrhenius and Włodzimierz Rabinowicz for making me aware of this possibility and Kagan's mention of it.

3. The Moral Insignificance of Interference

Consider again the cases 1 and 2. Assuming that the harm occurring when B is drowned by A is just as morally serious as the harm occurring when D drowns, **MMI** implies that A's drowning B is morally worse, *ceteris paribus*, than C's refusal to save D. This is a substantial moral claim, at odds with the claim that would typically be made by utilitarians. According to the latter, since the harm is as morally serious in both situations, both acts are as morally wrong, *ceteris paribus*.²¹

I will now support that the utilitarian view is superior to that of **MMI**. Since **MMI** is implied by any claim to the effect that the distinction is morally significant in itself, such an argument, if successful, shows, not only that **MMI** is untenable, but that the same goes for *any* ethical theory which claims the distinction to be of basic moral importance.

Consider the following illustration of case 1 and 2, designed to spell out the comparisons relevant for both **MMI** and utilitarianism.

Case 1

- | | |
|-------------------------|-------------------|
| (i) A drowns B | => B dies |
| (ii) A does not drown B | => B does not die |
| (iii) X is true of A | => B does not die |

Case 2

- | | |
|-------------------------------------|-------------------|
| (iv) C does not save the drowning D | => D dies |
| (v) C saves the drowning D | => D does not die |
| (vi) X is true of C | => D dies |

The acts whose degree of moral wrongness we are to compare are (i) and (iv). In order to do this from **MMI**, we only have to consider the result of (i) as compared to the result of (iii), and the result of (iv) as compared to the result of (vi). The results of (ii) and (v) may be disregarded. Utilitarianism, in contrast, compares the result of (i) with that of (ii) and the result of (iv) with that of (v), disregarding the results of (iii) and (vi). In other words, both **MMI** and utilitarianism consider some of the possible

²¹ This is not to say that a utilitarian cannot endorse the view that wrongness may come in degrees. See Eriksson 1994, chapter 6; and Eriksson 1997, for more about this. But since the moral judgement here is made *ceteris paribus*, such possibilities do not enter our considerations.

comparisons in each case to be irrelevant for the assessment of the moral wrongness of (i) as compared to that of (iv), but they differ in their view of *which* comparisons are irrelevant.

My point, now, is this: while the utilitarian can put forward an *independent argument* for disregarding the comparisons she disregards, the supporter of **MMI** cannot. This argument rests on a principle that I will call the *Principle of Attainability*, **PA** for short. In order to formulate **PA** in a clear manner, consider again the scenarios (i)-(vi). Each of these can be formulated as a conditional in which the antecedent describes a logically possible condition regarding the agent (for example, in (v) the antecedent describes the condition that C saves D). Let us call the conditions described by the antecedent of such conditionals *agent-conditions*, or, more formally put:

Z is an agent-condition =df. 1) Z is the antecedent of a conditional statement which describes what would be the upshot if certain conditions regarding an agent, *a*, were true. 2) Z describes a state of *a* such that it is logically possible for *a* to be in that state.

Now, **PA** can be formulated as follows:

PA: The fact that the truth of a certain agent-condition, *Z*, regarding an agent, *a*, in a situation, *s*, would effect a certain upshot, cannot make, or be a part of a fact that makes, *a*'s actual behaviour in *s* (more) morally wrong unless either (i) *Z* is in fact true in *s*, or (ii) *Z* is not true in *s*, but *a* in *s* have it in her power to see to it that *Z* is true in *s*.

Note that **PA** only sets limits for the moral appraisal of the *actual* behaviours of agents. Imagine that C cannot swim and that, therefore, if she were to try swimming out into the water to save D, she would not succeed. It is still true that *if* C were to swim out into the water (which she is in fact unable to do), D would be saved. **PA** does not contradict that *if* C would do this, the fact that D would then be saved is a relevant moral consideration for assessing this *hypothetical* behaviour of C (i.e., her swimming out into the water). What **PA** rules out is only that C's *actual* behaviour can be made (more) wrong by the fact that D would be saved if C were to swim out into the water, and the reason for this is simply that C in fact cannot swim out into the water.

It should also be noted that **PA** only sets limits to what may *make* actions

(more) wrong. **PA** does not rule out that a particular agent's *performance* of an action that is morally wrong is influenced or caused by facts about this agent which she herself cannot change, as long as these facts are not wrong-making characteristics in a basic ethical theory.²² For example, one thing that may influence that a certain agent acts wrong according to utilitarianism is the fact that she (like most people) is psychologically unable to perform and be guided by a complete utilitarian calculus. This influences her wrongdoing to the extent that if she had been such a "perfect utilitarian agent" she would have acted right according to utilitarianism. However, it is not *this* fact about her that *makes* her actual action wrong, but rather the fact that she does not perform the same action as she would have performed, had she been a "perfect utilitarian agent" – i.e., the fact that her actual action fails to maximise utility. This is not ruled out by **PA**.

The argument for disregarding the comparisons between (i) and (iii), and (iv) and (vi) now proceeds as follows.

1) (i) and (iv) contain agent-conditions describing A's and C's actual behaviour in two particular situations, *s* and *s'*.

2) (iii) and (vi) contain agent-conditions regarding A and C in *s* and *s'* respectively which are in fact not true.

3) A in *s* does not have in her power to make the agent-condition contained in (iii) true in *s*.

4) C in *s'* does not have in her power to make the agent-condition contained in (vi) true in *s'*.

5) **PA**

6) *Therefore*, the fact that certain upshots would result if the agent-conditions contained in (iii) and (vi) were true in *s* or *s'* cannot make the behaviours described by the agent-conditions contained in (i) or (iv) (more) morally wrong.

7) **MMI** implies that the fact that certain upshots would result if the agent-conditions contained in (iii) or (vi) were true in *s* or *s'* makes the behaviour described by the agent-condition in (i) more morally wrong than the behaviour described by the agent-condition in (iv).

8) *Therefore*, **MMI** is false.

More complex versions of **MMI**, where facts about interference are ascribed moral significance only when occurring together with other kinds of facts, can be attacked with an analogous argument. This argument will

²² I am grateful to Erik Carlson for pointing out to me this possible source of confusion.

be different only in that it will employ the part of **PA** ruling out that facts about which upshots that would follow from the truth of an agent-condition cannot even be a *part* of a fact that makes an action (more) wrong unless the conditions of **PA** are met.

The supporter of **MMI** is unable to use **PA** for a similar argument for disregarding the upshots that would result if the agent-conditions in (ii) and (v) were true. For the fact is that A in *s* has it in her power to see to it that the agent-condition in (ii) is true in *s*, and that C in *s'* has it in her power to see to it that the agent-condition in (v) is true in *s'*. If this would be denied, there would be no moral problem at all involved in (i) and (iv) – neither A nor C would be acting wrong to any degree. This follows from the principle that 'ought' implies 'can': if A and C did not have it in their power to see to it that the agent-conditions in (ii) and (v) respectively were true in *s* and *s'*, they would not have in their power to act as to avoid the deaths of B and D.

An important step in the argument is the premises 3 and 4 – the claim that A in *s* and C in *s'* does not have it in their power to see to it that the agent-conditions contained in (iii) and (vi) are true in *s* and *s'* respectively. If we return to the four suggested specifications of *X* presented in section 2 above, this is quite easy to see. No matter what C or A do, they cannot change the fact that they have ever existed, exist in *s* and *s'* respectively, are present in *s* and *s'* respectively, and are in a position where they can chose between alternatives which would make a difference to the harm suffered by B and D respectively.

Much of the strength of my argument against **MMI** thus comes down to the plausibility of **PA**. I see two reasons for accepting this principle. First, it is simply intuitively attractive and compelling. More important, however, its intuitive attraction can be theoretically explained in terms of **PA** being an extension of the principle that 'ought' implies 'can'. In the same way as the fact that a possible *action* would, if I performed it, have a certain upshot cannot make my actual behaviour (more) wrong unless that action either is what I actually do or is something I *can* do, the same can be said about *any* agent-condition regarding me. Since actions are just one sub-class of agent-conditions, **PA** is more basic than the principle that 'ought' implies 'can', and that the latter is supported by the former. In this way **PA** is not only supported by it being an extension of an intuitively very plausible principle, but also by its epistemological role to provide evidence for this principle.

The part of **PA** according to which facts about which upshots would

follow if certain agent-conditions were true cannot even be a *part* of a fact that makes an action (more) wrong unless they meet the conditions stated in this principle can be defended in a similar way. For it seems clear that the principle that 'ought' implies 'can' supports, not only the claim that facts about which upshots would be effected by the performance of various possible actions cannot make my actual behaviour (more) wrong unless I *can* perform those actions, but also that such facts cannot even be *parts* of more complex facts which make my actual behaviour (more) wrong unless this provision is met. Consider, for example, a possible version of utilitarianism claiming that what makes actions wrong is the complex fact that they effect a certain level of universal happiness and that there are other merely logically possible actions that would, if performed, have effected more universal happiness. This version of utilitarianism could be plausibly objected to by claiming that it violates the principle the 'ought' implies 'can', since some of the logically possible actions will not be such that an agent *can* perform them. The part of **PA** under consideration helps to support also this aspect of the principle that 'ought' implies 'can'.

Assuming that **PA** is sound, the supporter of **MMI** (or some more complex version of it) can try a number of moves in order to refute my argument. First, she could employ a different analytical framework than the counterfactual one hitherto used in the debate and show that such a framework is superior for analysing causal statements and that it rescues **MMI** from my argument. I concede this possibility, although I doubt its prospect for success. However, in light of such a move being successfully carried out it is up to the critics of **MMI** to formulate a *new* argument. My guess is that an argument structurally similar to the one stated here with the help of **PA** would be rather easily formulated also in a different framework for analysing causal statements. Secondly, she could try finding a description of the condition *X* in the antecedent of the counterfactual formula for the distinction which does not make **MMI** incompatible with **PA**. Again, I have my doubts whether this is a possible move, but concede that where it to be carried out successfully, the burden of proof again shifts to the critics of **MMI**.

Until any such arguments are presented, however, I conclude that the argument set forth above strongly supports the claim that the distinction between causing harm through interference and non-interference lack the kind of moral significance usually ascribed to it by its supporters. Facts in terms of this distinction cannot in themselves make, or be parts of other

facts that make, actions (more) morally wrong.²³

²³ Draft versions of this paper have been presented at research seminars at the department of philosophy, Göteborg University, and at *Filosofidagarna*, Swedish national conference in philosophy, Lund, 13-15 June, 1997. I thank the following people for comments and criticism: Gustaf Arrhenius, Lars Bergström, Erik Carlson, Patrik Christiansson, Jonas Gren, Johan Mårtensson, Włodzimierz Rabinowicz, Lars Sandman, Bolof Stridbeck, Jan Svensson, Ander Tolland, May Torseth and Torbjörn Tännsjö.

References

- BENNETT, J 1966, "Whatever the Consequences", *Analysis*, vol. 26, 1966, pp. 83-102
- 1981, "Morality and Consequences", in McMurrin, S (ed.), *The Tanner Lectures on Human Values*, ii, Salt Lake City 1981: University of Utah Press.
- 1995, *The Act Itself*, Oxford: Oxford University Press.
- DINELLO, D 1971, "On Killing and Letting Die", in Steinbeck, B & Norcross, A (eds.), *Killing and Letting Die*, New York 1994: Nordham University Press.
- DONAGAN, A 1977, *The Theory of Morality*, Chicago: University of Chicago Press.
- ERIKSSON, B 1994, *Heavy Duty. On the Demands of Consequentialism*, Edsbruk & Stockholm: Almqvist & Wiksell International.
- 1997, "Utilitarianism for Sinners", *American Philosophical Quarterly*, vol. 34, no. 2, 1997, pp. 213-228
- FINNIS, J 1973, "The Rights and Wrongs of Abortion: A Reply to Judith Jarvis Thomson", *Philosophy & Public Affairs*, vol. 2, pp. 117-45
- FOOT, P 1967, "Abortion and the Doctrine of the Double Effect", *The Oxford Review*, no. 5, 1967, pp. 5-15.
- 1985, "Morality, Action and Outcome", in Honderich, T (ed.), *Morality and Objectivity*, London 1985: Routledge & Keegan Paul.
- GAUTHIER, D 1986, *Morals by Agreement*, Oxford: Oxford University Press.
- GLOVER, J 1977, *Causing Death and Saving Lives*, Harmondsworth: Penguin Books.
- GOODMAN, N 1979, *Fact, Fiction and Forecast*, Cambridge Mass: Harvard University Press.
- GREENWELL, J R 1975, *The Morality of Abortion*, Ph.D. Diss., University of Arizona.
- HARE, R M 1981, *Moral Thinking: Its Methods, Levels and Point*, Oxford: Clarendon Press.
- JACKSON, F 1991 (ed.), *Conditionals*, Oxford, Oxford University Press
- KAGAN, S 1988, "The Additive Fallacy", *Ethics*, vol. 99, October, 1988, pp. 5-31.
- 1989, *The Limits of Morality*, Oxford: Oxford University Press.
- NOZICK, R 1974, *Anarchy, State and Utopia*, Oxford: Basil Blackwell.
- QUINN, W S 1989, "Actions, Intentions and Consequences: The Doctrine of Doing and Allowing", *The Philosophical Review*, vol. 98, no. 3, 1989, pp. 287-312.
- RACHELS, J 1975, "Active and Passive Euthanasia", *The New England Journal of Medicine*, vol. 292, 1975, pp. 78-80.