

DOING/ALLOWING AND THE DELIBERATIVE REQUIREMENT

Fiona Woollard

Abstract

Attempts to defend the moral significance of the distinction between doing and allowing harm directly have left many unconvinced. I give an indirect defence of the moral significance of the distinction between doing and allowing, focusing on the agent's duty to reason in a way that is responsive to possible harmful effects of their behaviour. Due to our cognitive limitations, we cannot be expected to take all harmful consequences of our behaviour into account. We are required to be responsive to harmful consequences that have some feature that makes it easy for us to become aware of them. I show that, under Jonathan Bennett's analysis of the doing/allowing distinction, harm that is incidentally done has such a feature, which is not shared by harm that is incidentally allowed. Any plausible analysis of the doing/allowing distinction will entail a similar asymmetry. It follows that, *prima facie*, an agent who incidentally does harm has violated a moral requirement (the deliberative requirement) which an agent who incidentally allows harm has not violated.¹

Many people believe that there is a morally significant distinction between doing harm to another and simply allowing that harm to occur. Other things being equal, it seems worse to do harm than to simply allow it to occur. Killing is worse than letting die; breaking Tom's leg is worse than allowing Tom leg to be broken.

However, attempts to argue directly that the distinction is morally significant have left many unconvinced. In this paper, I attempt to provide an indirect argument that this distinction is morally significant. Rather than looking directly at wrongness associated with the performance of doing or allowing harm, I focus on the agent's deliberative responsibilities, his obligation that his practical reasoning be appropriately responsive to potential harm to others. I argue that an agent who does harm has

¹ A substantial part of the research for this paper was part of a PhD funded by an AHRC Doctoral Studentship.

(prima facie) violated this deliberative requirement, whereas an agent who merely allows harm has (prima facie) not violated this requirement. This is enough to imply that there is a default moral distinction between cases of doing harm and cases of merely allowing harm.

This paper does not defend the full moral distinction between doing and allowing presupposed by commonsense morality. Nonetheless, its conclusion is still of interest. It shows that due to the nature of the distinction between doing and allowing, coupled with some fairly deep facts about human cognitive capacities, there is a morally significant asymmetry between cases of doing harm and cases of merely allowing harm. As the ambitions of my argument are relatively modest, it can be established using relatively uncontroversial premises and thus should be convincing to a wider audience. Such a modest but widely appealing argument is well worth having.

The Requirement to be Responsive and Its Limits

For the purposes of this paper, I assume that we are morally required to take the welfare of others into account in our practical deliberation: when another person may be harmed as a result of a contemplated action, we should become aware of this fact and see it as a reason against performing the action. It is not enough to simply avoid harming others; we need to deliberate in a way that is properly responsive to harm that could come to others as a result of our actions. For example, if I back my car out of the driveway without looking, failing to consider whether there may be any pedestrians in my path, I have failed in my obligations to reason responsibly. I call the obligation to reason responsibly the *deliberative requirement*.²

The deliberative requirement does not demand that agents consciously check each of the consequences of their actions, considering whether there is any potential harm to others. Instead the deliberative requirement demands alertness to potential harm, so that the agent is 'keeping his eyes open' for harmful

² This is a rough statement of the deliberative requirement. A person need not always become aware of potential harm to others; this demand is waived if, for example, more salient features of the situation already settle whether he should perform the action.

consequences. The agent's reasoning needs to be of such a quality that if he had reason to think an intended action would result in harm to others, he would become aware of this fact and see it as a reason not to perform the action.³

Commonsense suggests that we are required to be responsive to at least some potential harm to others. However, an agent cannot be required to be responsive to all foreseeable harm to others. The deliberative requirement must be limited. Therefore there must be cases in which (a) it is foreseeable that harm will occur if the agent acts in a particular way but (b) he does not violate the deliberative requirement if he fails to take this into account.

My argument appeals to considerations of difficulty, rather than considerations of cost.⁴ It is not within human cognitive capacities to be responsive to all potential harmful consequences of our acts while deliberating. This is because there are *too many* foreseeable consequences of each action for it to be humanly possible for us to be responsive to harm in all of them while deliberating. To attempt to be responsive to all foreseeable harm resulting from one's behaviour would be like attempting to pick out every word in common English with 'w' as its fourth letter. There may not be that many cases that I need to spot; I may know all the words in common English and, given any particular word, I know whether it has 'w' as its fourth letter. However, the pool of objects I am searching in is too large to be manageable. Even if I am not being careless, there is no way that I can expect to pick out all the relevant words. If we were required to undertake tasks like this before acting, we would never be able to act at all.

Without his reasoning being in any way flawed, an agent can fail to draw a conclusion that follows quite obviously from a set of accepted premises. This is because human agents cannot, and should not be expected to, survey all their beliefs at once and pick out all those which bear on the matter at hand. In practical

³ Kent Bach's 'default reasoning' model of theoretical and practical reasoning suggests that responsiveness to certain features of a situation is a common justificatory demand on our reasoning. See Kent Bach, 'Default Reasoning: Jumping to Conclusions and Knowing When to Think Twice,' *Pacific Philosophical Quarterly* 65 (1984), pp. 37–58.

See also Samuel Scheffler, *Human Morality* (New York, Oxford: Oxford University Press, 1992) pp. 29–51, especially p. 32, for a discussion of different ways moral considerations may impinge upon our deliberation that includes something very like responsiveness.

⁴ For the distinction between considerations of cost and difficulty see G.A. Cohen, *Karl Marx's Theory of History* (Oxford: Oxford University Press, 1978), p. 238 or G.A. Cohen, 'On the Currency of Egalitarian Justice', *Ethics*, 99 (1989), pp. 906–944.

deliberation, we face an additional difficulty: the pressure to deliberate quickly. Often, if we do not make a decision quickly, the opportunity will pass and thus we will have refrained from acting by default. If the deliberative requirement requires us to spend too long calculating the consequences of our actions, it will be paralysing.⁵ Plausible moral requirements cannot be paralysing in this way. It must be humanly possible for us to fulfil the deliberative requirement within a short enough time that action is possible.

Thus, considerations of difficulty prevent the deliberative requirement from being complete. This leads to the conclusion that the set of foreseeable incidental harms needs to be divided into harms that the agent is required to take into account in practical deliberation and harms that he is not required to take into account.

To pick out the vital set of consequences, I begin by considering the notion of salience. Salient facts 'leap out' at the deliberator: he finds them especially prominent or conspicuous.⁶ Salience is usually relative to a particular deliberator and a particular piece of reasoning: the facts that are salient to me in a given situation may not be salient to you. I require a more objective notion of salience, which can be applied to the deliberation of all (or almost all) human deliberators. This type of salience, referred to simply as salience henceforth, is still relative to a given piece of deliberation. However, it uses a standardised deliberator: a deliberator with normal human cognitive capacities who is deliberating with reasonable care and who has certain assumed concerns. A fact is salient in my sense if it would be salient to this standardised deliberator in a given piece of deliberation.

'Normal human cognitive capacities' may be either a normative notion or a statistical one: it may be the capacities such that an (adult) human without such capacities is a defective member of the species, or it may be the capacities that in fact most adult members of the species possess. We are interested in normal

⁵ There is a large literature discussing the counter-productivity of using act-consequentialism as a guide to action. This literature trades partly on the fact that it would be paralyzing to try to calculate which available action has the best consequences before acting. See Derek Parfit, *Reasons and Persons* (Oxford: Clarendon, 1984), pp. 24–9, 31–43; Peter Railton, 'Alienation, Consequentialism, and the Demands of Morality' *Philosophy and Public Affairs*, 13 (1984), pp. 140–6, 152–3.

⁶ Obviously, this is not an analysis of salience. It is not part of my project to analyse salience: indeed, there are serious questions about whether such an analysis can be given. I thank Maximilian de Gaynesford for helpful discussion of this point.

human capacities because we are interested in what is required of normal human agents. What is required of normal human agents will depend upon what is possible/difficult/impossible for normal human agents. Thus, even if the 'normal human' is a purely statistical notion, we can use the notion of normal human capacities to inform normative principles.

My standardised notion of salience uses the idea of 'reasonable care' in deliberation. Careless deliberators take into account only what is right in front of them; they ignore anything that is not immediately in their minds. Reasonable care requires more than this. Reasonable care involves an agent putting some effort into his deliberation, being responsive to some (but not necessarily all) relevant premises from his background belief system. Whether an agent has exercised reasonable care or not may remain a matter of judgement. However, we can recognise cases in which an agent has not exercised reasonable care in his deliberation; we can also pick out cases in which an unreasonable degree of care would be required for an agent to become aware of a given fact. We can therefore pick out facts that would be salient to a normal human agent using reasonable care without requiring a specific statement of what reasonable care consists in.

The standardised agent is assumed to have the concerns that are part of the piece of deliberation – what the particular agent wishes to achieve/avoid by the course of action under consideration. He is also assumed to have certain other concerns, including a concern that others should not be harmed as a result of his actions.

The deliberative requirement requires us to be responsive to precisely those facts about potential harm to others that are salient in our standardised sense. It is not a fault in our reasoning if we do not take into account significant facts that are not salient. This simply shows that our powers of reasoning are limited. However, it is a fault in our reasoning if we fail to take into account facts that are salient. A fact about potential harm to others is salient if it would be taken into account by an agent with normal human cognitive capacities, exercising a reasonable amount of care in his deliberation, who has a number of assumed concerns *including the concern that harm should not come to others as a result of his actions*. Thus if an agent fails to take a salient fact about potential harm into account then either (a) he does not have normal human cognitive capacities; (b) he has not exercised reasonable care in his deliberation or (c) he is not concerned to ensure that others do not suffer harm as a result of his actions. If

(a) is true, and the agent is not culpable for this fact, then he may be excused from any wrongdoing resulting from his incapacities. If (b) or (c) is true, then the agent is guilty of precisely the type of moral flaw to which the deliberative requirement speaks.

Failure to take a fact about potential harm to others will be a violation of the deliberative requirement just in case that fact is salient. What makes a fact salient in a given piece of deliberation? One feature that makes a fact salient is its significance – its bearing on the matter at hand. A fact is significant in a given piece of deliberation, if, given the assumed set of concerns of the standardised deliberator, if he knew that fact was true, he would see this as relevant to his deliberation. However, significance alone is not enough to make a fact salient. A fact can be very significant and yet fail to be salient. Some facts (for example, the fact that a certain method would cure cancer) are important, but are very difficult to work out even if they follow from things we already know.

If the significance of a fact is not enough to make it salient, something else must be necessary. There must be some other features which, combined with a fact's significance, make it salient. These are features of the fact that make it easier for us to become aware of it, should we think it significant. There is, we might say, a match between these features and our natural patterns of deliberation; these features characterise the types of fact that human cognitive capacities are naturally fitted to pick up on. I will refer to these features as *non-evaluative salient making features* because they combine with the significance of a fact (how the deliberator evaluates its importance) to make it salient.

Facts with non-evaluative salient making features are not necessarily those that we have most reason to consider. Instead, they are the type of facts that we find it easiest to become aware of. For analogy, suppose that in a game I am asked to name all the red objects in the room I have just left. I will pick out the largest red objects in a room first, remembering a red chair before a red pencil-sharpener. Given very limited observation time, I will probably pick out only the largest red objects. Size is something that makes an object particularly obvious to us – our natural capacities for observation make us more likely to spot large objects than small ones. Similarly, a fact with non-evaluative salient making features is a fact that our natural cognitive capacities make us particularly likely to 'spot' when we are reasoning.

The fact that an action is likely to result in harm to others is a significant fact in deliberation about whether to perform that

action. Non-evaluative salient making features make significant facts salient. Thus if a fact about potential harm to others has a non-evaluative salient making feature, it will be *prima facie* salient. The deliberative requirement requires us to be responsive to salient facts about potential harm. So we are *prima facie* required to be responsive to any fact about potential harm that has a non-evaluative salient making feature.

Non-evaluative Salient Making Features and the Doing/Allowing Distinction

To complete my argument, I need to show that harm the agent would do has a non-evaluative salient making feature that is not present in harm he would merely allow. If this is so, then *prima facie* harm the agent would do is salient while harm he would merely allow is not. It follows from the above argument that we are morally required to be responsive to harm we do but we are not morally required to be responsive to harm we allow. There is a moral asymmetry between doing and allowing.

I will show that such an argument can be given in cases of incidental harm i.e. cases where the harm is a foreseeable by-product of the agent's pursuit of an otherwise legitimate goal, but it is neither a goal of the agent, nor a means to any of his goals. When the harm is either one of the agent's goals or a means to his goal, it will clearly be salient in his deliberation. Thus the most we can expect from the deliberative requirement is an asymmetry between doing and allowing *incidental* harm.

I do not want to commit myself to any particular analysis of the doing/allowing distinction here. I believe that any plausible analysis will leave doings with a non-evaluative salient making feature not present in allowings. However, in this paper, I will focus on Jonathan Bennett's analysis of the doing/allowing distinction. I focus on Bennett's account for two reasons. First, it is one of the best analyses of the doing/allowing distinction in the literature. Second, it is widely agreed to pick out a distinction that has no moral relevance.⁷ It is in no way obvious that Bennett's analysis has

⁷ Jonathan Bennett, *The Act Itself* (Oxford: Oxford University Press, 1995), p. 105; Jeff McMahan, 'A Challenge to Common Sense Morality', *Ethics*, 108 (1998), p. 397; Judith Jarvis Thomson, 'The Act Itself: Review', *Nous* 30 (1996), p. 550.

a connection to salience. Thus this account presents a particular challenge for my argument. However, my argument does not depend on acceptance of Bennett’s analysis of the doing/allowing distinction, for other accounts can also be shown to entail that doings possess a non-evaluative salient making feature not present in allowings.

Bennett analyses our notions of doing and allowing in terms of the distinction between positive and negative relevance to an outcome. He distinguishes between positive and negative relevance to an outcome in terms of the proportion of ways the agent could have moved his body without that outcome occurring. Roughly speaking, an agent is positively relevant to an outcome if only a small proportion of the ways he could act would lead to the specified outcome; he is negatively relevant if a very large proportion of the ways he could act would lead to the specified outcome.

Bennett uses a square, called the agent’s behaviour space, to represent the agent’s possible movements. If an agent is negatively relevant to an outcome then the proposition ‘Outcome occurs’ will correspond to a large subspace of the behaviour space. This subspace will be much larger than the subspace corresponding to the proposition ‘Outcome does not occur.’ The converse will be true if the agent is positively relevant. This is illustrated in the following diagrams:

Figure 1 Negatively Relevant to Outcome





Figure 2 Positively Relevant to U



My argument requires that I pick out an non-evaluative salient making feature that, under Bennett's analysis, attaches to incidental doings but not incidental allowings. Consider the following pair of subjunctive conditionals:

Would-if-I-Did Conditional: If the agent were to adopt the planned course of behaviour, the outcome would occur.

Wouldn't-if-I-Didn't Conditional: If the agent were not to adopt the planned course of behaviour, the outcome would not occur.

Suppose an agent is deciding whether to follow a certain course of behaviour and a given outcome would occur incidentally if he does so. I claim that, if it were foreseeable to an agent that he fulfils both the Would-if-I-Did Conditional and the Wouldn't-if-I-Didn't Conditional with respect to the outcome, this would make the fact that the outcome would occur non-evaluatively salient. In other words, the fact that the outcome may occur incidentally will be non-evaluatively salient if it is foreseeable that the outcome would occur if he were to adopt the planned course of behaviour, but would not occur if he were not to adopt that course of behaviour.

So I need to show (1) that fulfilling both the Would-if-I-Did Conditional and the Wouldn't-if-I-Didn't Conditional is an non-evaluative salient making feature and (2) that Bennett's analysis implies that outcomes incidentally brought about fulfil both the Would-if-I-Did Conditional and the Wouldn't-if-I-Didn't Conditional but that outcomes incidentally allowed do not do so.

The argument for (2)


Suppose an agent is considering a course of behaviour which would lead to his being incidentally positively relevant to a harmful outcome. (He would be positively relevant to the outcome, but the outcome's occurrence would be neither an end nor a means for him.) Consider:

Incidental Push: A vehicle stands unbraked on ground that slopes down to a cliff top. Agent intends to run quickly along a path as part of his fitness training. Due to the vehicle's position

on the path, if Agent runs past, he will knock the vehicle. It will then roll over the cliff edge to its destruction.⁸

In Incidental Push most of the ways that Agent could move would not lead to the vehicle's destruction. So by Bennett's analysis, if Agent runs along the path he will be positively relevant to the vehicle's destruction. Agent's behaviour space is divided in the following way:

Figure 3 Incidental Push



In Incidental Push, the following pair of subjunctive conditionals holds: If Agent were to run past quickly, the vehicle would be destroyed; if Agent were not to run past quickly, the vehicle would not be destroyed. Of course, there are a fair number of ways the vehicle might be destroyed without Agent running past quickly: for example the vehicle would be destroyed if Agent hopped to the vehicle and pushed it. However, my claim that the pair of conditionals holds is not undermined by these alternative ways of destroying the vehicle. In the great majority of cases in which the agent does not run past, the vehicle is not destroyed.⁹

Recall that an agent fulfils the Would-if-I-Did Conditional and the Wouldn't-if-I-Didn't Conditional with respect to an outcome if and only if:

Would-if-I-Did Conditional: If the agent were to adopt the planned course of behaviour, the outcome would occur.

Wouldn't-if-I-Didn't Conditional: If the agent were not to adopt the planned course of behaviour, the outcome would not occur.


⁸ This is a version of Bennett's *Push* example. See Bennett, *The Act Itself*, p. 67. Realistically, there will almost always be some ways of running that do not knock the vehicle hard enough to send it down the hill. I ignore such complications here as they do not affect the main thrust of the argument.

⁹ I assume that nothing else will lead to the vehicle's destruction if Agent's conduct does not.

So Agent fulfils both the Would-if-I-Did Conditional and the Wouldn't-if-I-Didn't Conditional with respect to the vehicle's destruction in Incidental Push.

To apply this to the general case, suppose that Agent will be positively relevant to the outcome if he follows a planned course of behaviour. The behaviour space is divided as shown below.

Figure 4 Planned course of behaviour is positively relevant to outcome



The following pair of subjunctive conditionals holds: If Agent were to perform the planned course of behaviour, the outcome would occur (Would-if-I-Did Conditional); if Agent were not to perform this course of behaviour, the outcome would not occur (Wouldn't-if-I-Didn't Conditional). So Agent fulfils both the Would-if-I-Did Conditional and the Wouldn't-if-I-Didn't Conditional with respect to this outcome.

The truth of both these subjunctive conditionals follows from Bennett's definition of positive relevance. The first conditional, that if Agent were to perform the course of behaviour, the outcome would occur, is obvious. The second conditional is slightly less obvious. As the outcome will occur if the agent performs the planned course of behaviour, then if the outcome does not occur, the agent cannot have performed the planned course of behaviour. So the area corresponding to 'Outcome does *not* occur' is part of 'Agent does *not* perform the planned course of behaviour'. The area corresponding to 'Agent does not perform the planned course of behaviour' is divided between 'Outcome does not occur' and 'Outcome occurs *and* Agent does not perform the planned course of behaviour'. By definition, if the Agent would be positively relevant to the outcome, then 'Outcome does not occur' takes up most of the behaviour space'. The area corresponding to 'Outcome does not occur' is much bigger than the area corresponding to 'Outcome does occur', which is obviously bigger than the area corresponding to


‘Outcome occurs *and* Agent does not perform the planned course of behaviour’. So the area corresponding to ‘Outcome does not occur’ is much bigger than the area corresponding to ‘Outcome occurs *and* Agent does not perform the planned course of behaviour’. So ‘Outcome does not occur’ takes up most of ‘Agent does not perform the planned course of behaviour’. So most of the ways in which Agent could act without performing the planned course of behaviour would not lead to the outcome occurring.¹⁰

We shall now consider cases in which if the agent were to follow a planned course of behaviour he would be negatively relevant to a harmful outcome. Consider:

Incidental Stayback: Another vehicle is already rolling down the slope a small distance away. Agent could interpose a rock to stop the vehicle rolling. Agent intends to continue running along the path. If he does so, he will not interpose the rock and the vehicle will roll to its destruction.¹¹

In this case, we have the following division of the behaviour space:

Figure 5 Incidental Stayback



Consider the conditionals we discussed above. The analogous conditionals in this situation are:

(Would-if-I-Did Conditional) (Stayback) If Agent runs along the path, the vehicle will be destroyed.

(Wouldn't-if-I-Didn't Conditional) (Stayback) If Agent does not run along the path, the vehicle will not be destroyed.


¹⁰ Formally, if B = Agent follows planned course of behaviour and O = outcome occurs: $B \subset O$ so $\neg O \subset \neg B$. $\neg B = \neg O \cup (O \cap \neg B)$; $\neg O > O > (O \cap \neg B)$; So most of $\neg B$ is taken up by $\neg O$.

¹¹ This is an elaboration of Bennett’s Stayback. See Bennett, *The Act Itself*, p. 67.

The Would-if-I-Did Conditional holds in Stayback. However, Wouldn't-if-I-Didn't Conditional does not hold. It is not the case that if Agent does not run along the path, the vehicle will not be destroyed. Most of the other ways Agent could move would still result in the vehicle being destroyed.

The same is true in the general case. Suppose that Agent would be incidentally negatively relevant to a harmful outcome if he were to perform the planned course of behaviour. Then the behaviour space is divided as shown below:

Figure 6 Planned course of behaviour is negatively relevant to outcome



The relevant conditionals are:

(Would-if-I-Did Conditional) If Agent were to adopt the planned course of behaviour, the outcome would occur.

(Wouldn't-if-I-Didn't Conditional) If Agent were not to adopt the planned course of behaviour, the outcome would not occur.

The Would-if-I-Did Conditional still holds. However, the Wouldn't-if-I-Didn't Conditional does not hold. It is not the case that if the agent did not perform the planned course of behaviour, the outcome would not occur. Most ways that the agent could move without performing the planned course of behaviour would result in the outcome occurring. Again, this follows from Bennett's definition of negative relevance to an event.

Therefore, on Bennett's analysis, outcomes that are incidentally done fulfil both the Would-if-I-Did Conditional and the Wouldn't-if-I-Didn't Conditional, whereas outcomes that are incidentally allowed do not fulfil both the Would-if-I-Did Conditional and the Wouldn't-if-I-Didn't Conditional.

The Argument for (1)

We now need to show that fulfilling both the Would-if-I-Did Conditional and the Wouldn't-if-I-Didn't Conditional is a non-evaluative salient making feature. This premise claims that the fact that an outcome may occur incidentally will be non-evaluatively salient if it is foreseeable to the agent that the outcome would occur if he were to adopt the course of behaviour under consideration, but would not occur if he were not to adopt this course of behaviour.

According to this premise, when the Agent in Incidental Push considers whether or not to run along the cliff path to get fit, the fact that in doing so he would knock the vehicle over the edge will be non-evaluatively salient. This fact is non-evaluatively salient because the vehicle would roll over the cliff if Agent were to run past it, but would not roll over the cliff if Agent were not to run past it.

Premise 1 picks out a feature of our natural patterns of deliberation. To make this clearer, let us consider an example of a decision-making process. I use an example where prudential reasons prevent an action so that our thoughts are not clouded by moral considerations.

Driving to Campus: I need to go to campus. A method of achieving this goal occurs to me: I could drive there. However, if I travel by car, I will use expensive petrol. If I don't travel by car, then I won't use petrol. I decide to walk.

In my decision-making process in *Driving to Campus*, I consider a consequence of my behaviour that would occur if I did follow a particular course of action, but would not occur if I did not follow that course of action.

There are various ways in which I could use the same amount of petrol without driving to campus. I could siphon the petrol off and pour it on the ground. But this does not alter the fact that I would use up the petrol if I drove to campus but would not use up the petrol if I did not drive to campus. Out of all the possible ways I could act, only a very small proportion involve using up the petrol. We can still say that the petrol would not be used if I did not drive to campus.

In the *Driving to Campus* example, I would use the petrol if I drove to campus, but would not use it if I did not drive to campus. I fulfil the subjunctive conditionals, the Would-if-I-Did Conditional and the Wouldn't-if-I-Didn't Conditional, with respect to the event *the petrol is used up*. Clearly, the fact that if I drive I will use petrol is non-evaluatively salient. If someone in a similar situation did not take this fact into account, we would conclude either that he was reasoning carelessly or that he did not really mind how much petrol he used up. Therefore we have a case in which the fact an event may occur incidentally is non-evaluatively salient in an agent's deliberation *and* the agent fulfils both the Would-if-I-Did Conditional and the Wouldn't-if-I-Didn't Conditional with respect to this fact.

However, the fact that I will use the petrol may be salient for some other reason, rather than *because* I fulfil the two conditionals with respect to it. There seems to be one strong competing hypothesis. This hypothesis claims that in deciding whether to drive to campus, I think of a set of alternative ways of achieving my goal (I can drive or I can walk) and then weigh up the good and bad outcomes associated with doing one rather than the other. According to this hypothesis, the non-evaluatively salient events are not those that would occur if I drove to campus but would not occur if I did not drive to campus. Instead, the non-evaluatively salient events are those that would occur if I drove to campus but would not occur if I walked to campus.

Often, however, we decisively reject a course of behaviour without having any alternative in mind. For example, Bob decides not to go to London on Saturday because it is too noisy; he need not know what he is going to do instead. A sufficiently undesirable feature is enough to put us off a plan of action even if we have no alternative plan. I claim that often the undesirable features that we invoke are those that would occur if we did adopt the course of action but would not occur if we did not.

In any case, my argument can survive a slight modification in response to the alternative hypothesis. Suppose that in our deliberation what we really do is compare the costs of a course of behaviour relative to certain alternatives. In order to deduce Premise 1 from this, all we require is the further premise that *not adopting the course of behaviour* must always be one of the alternatives that we consider. However precious a goal is to us, we should

surely be prepared to abandon it if pursuing it would result in sufficiently bad harms to others. There *may* be a select few goals that are not subject to this constraint, purposes that a person must be allowed to pursue come what may, but most everyday plans do not fall into this category.

The above discussion suggests that Premise 1 reflects the natural structure of our deliberation.¹² When we are deliberating about whether to adopt a particular course of behaviour we consider incidental outcomes that would occur if we were to adopt that course of behaviour, but would not occur if we were not to adopt that course of behaviour. Fulfilling both the Would-if-I-Did Conditional and the Wouldn't-if-I-Didn't Conditional is a non-evaluative salient making feature.

In addition, I showed above that on Bennett's analysis outcomes that are incidentally done fulfil both the Would-if-I-Did Conditional and the Wouldn't-if-I-Didn't Conditional, whereas outcomes that are incidentally allowed do not fulfil both these conditionals. Thus Bennett's analysis entails that harm we incidentally do has a non-evaluative salient making feature that harm we incidentally allow does not have. So incidentally doing harm will violate the deliberate requirement while incidentally allowing harm will not usually do so.

The Doctrine of Doing and Allowing

When we deliberate about what to do, we are morally required to be responsive to potential harm to others. However, we cannot be required to be responsive to all harmful potential consequences, for we are not cognitively capable of doing so. We (at least those of us with normal human cognitive capacities) are required to be responsive to harmful consequences that are salient relative to a standardised deliberator. I have shown Jonathan Bennett's positive/negative relevance analysis, which is widely assumed to imply that the distinction is morally neutral, entails that harm incidentally done has a feature that harm that is incidentally allowed does not have. I have suggested that this feature is a

¹² Of course I have not proved that fulfilling Conditional 1 and Conditional 2 is a non-evaluative salient making feature. I do not think a philosopher could offer such a proof. An experimental psychologist might be able to do so.

non-evaluative salience-making feature. I suggest that any plausible account of the doing /allowing distinction will entail that doings have a non-evaluative salient making feature not associated with allowings. As any fact about potential harm to others with a non-evaluative salient making feature is *prima facie* salient, it follows that we are *prima facie* required to be responsive to incidental harm that we do, but not to be responsive to incidental harm that we allow. Thus an agent who does harm incidentally has *prima facie* violated a moral requirement that an agent who allows harm incidentally has *prima facie* not violated.

This argument supports a limited moral distinction between doing and allowing. My conclusion is limited in three ways. First, it applies only to harm that is done (or allowed) incidentally. As noted above, harm is incidental if it is a foreseeable by-product of the agent's pursuit of an otherwise legitimate goal, but it is neither a goal of the agent, nor a means to any of his goals.

Second, I argue only that *prima facie* an agent who has incidentally done harm will have violated a moral requirement that the agent who has incidentally allowed harm will not have violated. In some cases an agent who allows harm does violate this requirement; in other cases an agent does harm but does not violate the requirement. However, in these cases, there is always some relevant consideration that changes the default position. This happens if, for example, the harm incidentally allowed by an agent is immediately present to him – i.e. if a child lay bleeding beside his path.

Third, I do not argue here that the act of doing harm is itself worse than the act of allowing harm. I argue only that *prima facie* an agent who incidentally does harm violates a moral requirement that an agent who incidentally allows harm does not violate.

Despite these limitations, the conclusion is still of interest, for it shows that doing and allowing are not morally equivalent. It shows that the doing/allowing distinction is morally significant in the following sense: an agent who does harm can have violated a moral requirement that an agent who allows harm has not violated even if all other relevant factors are *apparently* equal i.e. neither agent has malicious intentions, the cost to each agent of acting differently is the same, etc. Moreover, because my argument supports a *prima facie* moral difference, it means that the moral asymmetry is the default position. For cases of doing and allowing to be morally equivalent there will need to be some unusual feature of the situation that cancels the moral

significance of the distinction. Such features will indeed be unusual. Most cases of doing and allowing harm are incidental by-products of our pursuit of other goals. Most cases of allowing harm do not involve non-evaluative salient making features such as the immediate presence of the victim. Thus to demonstrate the existence of this default distinction is enough to show that there is an important moral difference between doing and allowing.

*Department of Philosophy, University of Sheffield
45 Victoria Street, Sheffield, S3 7QB
f.woollard@sheffield.ac.uk*