

Intentionality, Intending and Moral Responsibility

Wesley H. Bronson
Princeton University

Many philosophers point to the so-called “doctrine of double effect” to show the tension between intending, intentional action and moral responsibility. The classic example goes as follows: a bomber has been given the assignment to bomb a munitions plant which is producing weapons used against his country, so he must go and bomb the plant. In bombing, however, he will undoubtedly kill innocent civilians in the nearby area of the blast. In this case, the bomber intends to destroy the plant (he desires its destruction), and he intentionally bombs it (he brings about the known consequence of his action). But while he also intentionally kills innocent civilians, he may not intend to kill them. The debate is thus whether or not the bomber is morally responsible for the death of the innocent civilians whom he intentionally kills but doesn’t intend to. With this scenario in mind, I argue in this paper that simply the fact that a consequence is brought about intentionally is insufficient for an agent to be held morally responsible, and further, that intending and being causally responsible for a consequence are often necessary to be morally responsible.

Wesley H. Bronson

Many argue that an agent intentionally brings about any known consequence of the action he performs. So if an action will result in consequences C_1 , C_2 and C_3 , and the agent knows that a particular action will result in all three consequences, then in performing the action, the agent ostensibly intentionally brings about those consequences. There are two advantages to defining action in such a way. First, it is very simple. If an agent chooses to perform an action, then any known consequence is intentionally brought about. Second, issues over moral responsibility can plausibly be settled easily. An agent can plausibly be held morally responsible for bringing about any known consequence.¹ After all, he knew the consequence would result, and he performed an action that would bring it about. So perhaps an agent ought to be responsible for anything brought about intentionally.

Let's say I plan on burning down an apartment building. And let's say that I know that there will be various consequences of burning it down. I know many people in the building will die. I also know that the flames from the building will cause nearby structures to be somewhat damaged. And finally, I know that the fire department will spend a lot of time and money to put out the fire, probably even causing injury to several fire fighters. Given that I know ahead of time that my action (burning down the apartment building) will result in several different consequences (deaths, structural damage, etc.) it can be said that I intentionally bring about all those consequences. The simple view of intentional action mentioned above might simply say that since the consequences were intentionally brought about, I am morally responsible for all of them.

This position, however, is overly simple, and fails to take into account the mental state of the agent. Instead, the position I will use in this paper, therefore, is similar to the more complex position accepted by Michael Bratman. Rather than group all types of action for which there is a known consequence to be an intentional action, Bratman distinguishes between two types. There is a difference between *intending* a consequence and *intentionally* bringing about that consequence.² Like the previous description, any action which brings about a known consequence is performed intentionally. But what if I must bring

1 Bentham, Jeremy. *An Introduction to the Principles of Morals and Legislation*. Ed. J Burns and H. Hart. London: Methuen, 1970, p.84.

2 Bratman, Michael. *Intention, Plans, and Practical Reason*. Cambridge: Harvard University Press, 1987. Chapter 10.

Intentionality, Intending and Moral Responsibility

about a known consequence as a side-effect of bringing about another consequence? What if I know a consequence will result, but I do not want that consequence, yet perform the action anyway? The simple view of intentional action will say that because it is a known consequence, I want to bring about that consequence. But what can be said of my mental state in which I simply do not want a certain consequence to result but know it will inevitably happen if I perform the action? In that sense, I *intend* to bring about only one of the consequences, but I nevertheless *intentionally* bring about all of them. In short, if an agent wants to perform a given action A because he wants consequence C₁, but he knows or is justified in believing that doing A will also result in consequences C₂ and C₃, then in doing A the agent can be said to intend C₁ but intentionally cause C₁, C₂ and C₃.

Many philosophers point to the so-called “doctrine of double effect” to show the tension between intending, intentional action and moral responsibility. The classic example goes as follows: a bomber has been given the assignment to bomb a munitions plant which is producing weapons used against his country, so he must go and bomb the plant. In bombing, however, he will undoubtedly kill innocent civilians in the nearby area of the blast. In this case the bomber intends to destroy the plant and he intentionally bombs it. But while he also intentionally kills innocent civilians, he may not intend to kill them. The debate is thus whether or not the bomber is morally responsible for the death of the innocent civilians whom he intentionally kills but doesn’t intend to. With this scenario in mind, I will argue in the remainder of this paper that simply the fact that a consequence is brought about intentionally is insufficient for an agent to be held morally responsible, and further, that intending and being causally responsible for a consequence are necessary to be morally responsible.

There is one exception which must be pointed out. Even those who would establish that any consequence performed intentionally results in the agent being held morally responsible would concede that an agent can be held morally responsible for actions that are not performed intentionally. That is, intentionality may be a sufficient but not a necessary condition for moral responsibility. Actions may be performed unintentionally through negligence for which we are still responsible. In cases of negligence, an agent can—but not necessarily— be held responsible for a consequence that he neither intended nor intentionally brought about. Take a drunk driver, for example. If he carelessly, in his

Wesley H. Bronson

drunken state, runs over and kills a pedestrian, few would deny that he ought to be held morally responsible for the death of the pedestrian.³ Note, however, that the driver did not intend for the death to occur, nor did he intentionally run over the pedestrian. In the case of negligence, he merely is causally responsible, yet nevertheless remains morally responsible. That is not to say that an agent ought to be held morally responsible in all cases of neglect. If I flip a switch that looks just like a light switch, thinking it to be a light switch, and instead it turns out to be an alarm, how can I be held blameworthy for flipping the alarm? I may be causally responsible for doing so, but since I had no possible way to know that the alarm would go off, I cannot be held morally responsible for the action.⁴ So while there are indeed times when a negligent action can result in the agent being held morally responsible, it need not be the case if, for example, there was no possible way the agent could have known the consequences that he ought to be held responsible there as well.

Similarly, Carl Ginet claims that there are indeed cases when an agent both *should* have known certain information and can then be held responsible for not knowing it, even if the consequence was brought about unintentionally. If Simon went to flip a light switch, but instead of looking like a normal switch it looked identical to a fire alarm, and upon flipping the switch the fire alarm went off, it seems rational, Ginet claims, to hold him morally responsible for making the mistake.⁵ There are certain times when evidence ought to be compelling enough that the information—in this case that the switch was, in fact, a fire alarm and not a light switch—should simply become apparent. Under these situations, even if the individual remains oblivious to the situation, he can still be held responsible. As Ginet concedes, “a certain amount of indignation towards him, for his causing the alarm to go off, would be deserved (though, of course, not as much as if he had intentionally set it off).”⁶ That is, we may be less angry at the agent for his lack of intention, but he is ostensibly no less blameworthy than an evil agent. Cases of negligence, therefore, when there is a lack of practical reflection on the action and its consequences admittedly seem to undermine my thesis.

3 Mele, Alfred and Steven Sverdlik. “Intention, Intentional Action, and Moral Responsibility.” *Philosophical Studies* Vol. 82, (1996). p. 269.

4 Ginet, Carl. “The Epistemic Requirements for Moral Responsibility.” *Philosophical Perspectives 14: Action and Freedom*, 2000. Ed. James E. Tomberlin. Boston: Blackwell Publishers, 2000. p. 269.

5 Ibid., 271.

6 Ibid., 271.

Intentionality, Intending and Moral Responsibility

My aim, therefore, must be somewhat clarified. Earlier I said that my aim was to show that establishing that an action was performed intentionally is insufficient to establish moral responsibility and that intending is a necessary requirement for moral responsibility. In light of cases of negligence, the thesis must be somewhat qualified, and I admit that maintaining intending as a necessary condition for moral responsibility only holds in cases where the agent reflects on his actions before doing it.⁷

Let us finally turn to intentional action. Why in the double-effect bomber example would someone be inclined to hold the bomber morally responsible for the death of the innocent civilians? Perhaps it is because the action was performed intentionally. Initially, the argument seems intuitively plausible. If the bomber could foresee the consequences and knew that the children would be killed, and he nevertheless bombed the munitions plant, then the fact that the action was performed intentionally would make him responsible for the deaths. But under more careful scrutiny, the fact that the action was performed intentionally is insufficient to establish moral responsibility. Let us look at a simple case of coercion to show that this is the case. Suppose Smith threatens to burn down Bob's house unless he throws a rock through a nearby window. Bob loves the building, however, and strongly opposes breaking the window. Stampeded by the threat, Bob unfortunately throws the rock, and the window breaks. Can he be held responsible for breaking the window? It seems that he intentionally breaks the window, but he is not morally responsible for doing so. Establishing that an action is performed intentionally is therefore not sufficient to establish moral responsibility.

There is a problem with the above example, however, and an obvious objection arises. Although in the case provided the individual did act intentionally and he can't be held morally responsible, he had no alternative possibilities. We tend to think that if an agent has a choice of whether or not to perform a given action, he can be held responsible for his choice. Carl Ginet comments, for example, that an individual must at some point have a choice of whether or not to perform a given action if we are to hold him morally responsible. As he defines a preliminary

⁷ My purpose here is to exclude cases of negligence. Negligence represents a unified and widespread class of phenomena. In a case of negligence, an agent can reflect on his action before doing it, but for some reason fails to do so. So we cannot say an agent must always intend a given consequence to be held morally responsible. My aim is to show now that in cases where the agent *does* reflect, then intending is a necessary condition for moral responsibility.

Wesley H. Bronson

definition of the *could-have-done-otherwise* (CDO) condition: “until t_1 it was open to S not to make movement M then or any other movement that would bring about H.”⁸ That is, at a given time (t_1), an individual (S) must have a choice of whether or not to perform a certain movement (M) that will bring about a certain consequence (H) if we are to hold him responsible. In the Bob and Smith example, Bob (S) does not have the choice of whether or not to throw the rock (M). He is stampeded by the threat, so he could not have done otherwise. Those who hold that moral responsibility requires that the agent intentionally bring about that consequence will say that the CDO condition was not met in the coercion example. Since he had no alternatives, of course he can’t be held morally responsible!⁹

As an interesting note, under similar conditions, we might ask what else *would* be required for Bob to be held morally responsible. Let’s say that again, Bob is threatened by Smith, but this time, Bob hates the building and wants nothing more than to cause damage to it. He willingly throws the rock, and the window breaks. Is Bob morally responsible for breaking the window in this second case? It seems we are inclined to think yes, he is responsible. Despite the threat, Bob intended to break the window and acted intentionally. The only difference between the two cases is that in the first, Bob did not intend to break the window, but in the second case, he did intend the consequence. We might be inclined to say that intending is a necessary condition in this case for establishing moral responsibility.¹⁰

The objection over alternate possibilities is legitimate, so let us therefore turn to a case where the agent intentionally performs an action where he had alternate possibilities, yet is still not morally responsible for the consequence. Such conditions could plausibly satisfy the previous objection. Mele and Sverdlik consider a doctor who must cause a patient some degree of pain during a procedure. The doctor does everything he can to minimize the pain, and he wishes there could be no pain at all, but it is an inevitable consequence of the procedure. They claim that

8 Ibid., 268.

9 See also, Frankfurt, Harry G. “Alternative Possibilities and Moral Responsibility.” *The Journal of Philosophy*. Vol. 66, No. 23, (Dec 4, 1969), pp. 829-839.

10 As Frankfurt similarly notes, if the individual were stampeded by the threat, we would not hold him responsible for the consequences of the action. But if the threat plays no role in leading the individual to action then he performs the action of his own volition. In this case, while he had no alternate possibilities, we can nonetheless hold him morally responsible.

Intentionality, Intending and Moral Responsibility

“insofar as an agent who is A-ing is neither aiming at A-ing nor trying to A, either as an end or as a means to (or constituent of) an end, she is not intentionally A-ing.”¹¹ Mele defends his dentist by claiming he did not specifically aim at causing the pain, and the action was thus non-intentional.¹² But what if an agent does aim at causing pain? Can his actions still be vindicated in some way?

As I will show, it is entirely possible for an agent to aim at causing pain—thus intentionally bringing about that specific consequence—yet not be held morally responsible for that consequence. At first, this statement may sound intuitively implausible, so an example is necessary. Andrew is standing with his friend Bob, both of whom have lived moral lives. Because of unfortunate circumstances, however, Bob is in a position where unless someone intervenes, he will die. Let’s say that a wild, poisonous snake is about to bite him. Andrew recognizes that the snake is about to bite Bob, and sees no way to save him except by stomping on his foot, thus causing Andrew pain which will cause him to get out of the way of the snake.¹³ Causing Bob this small pain seems to be the only way to make Bob get out of the way of the snake. In such a situation, Andrew intentionally brings about Bob’s pain by stomping on his foot. And Mele and Sverdlik would agree that the action was intentional, for Andrew aimed at causing Bob pain. Nevertheless, would we really hold Andrew morally responsible for the pain Bob experienced? After all, he saved Bob’s life by doing so.

There are now two ways in which we may analyze Andrew’s action. On the one hand, we may say despite the consequence he is still morally responsible. On the other hand, we may say he is not morally responsible for either choice he makes, including causing Bob’s pain. Let’s first look at the former situation. It is possible that Andrew is, in fact, morally

11 Ibid., 274. (In performing “intentional action,” Mele and Sverdlik imply that the individual does not intend the consequence. My aim here is to show that the middle ground of non-intentional action does not vindicate the agent in any way.

12 Two ideas must be clarified here. First, they describe intentional action in such a way that the agent does something intentionally when he intends the consequence. Second, the authors vindicate moral responsibility when the agent does not aim at causing the patient pain, or when the pain is a means to a better end. I will discuss this latter consideration later, but for now we are not concerned with consequentialist reasoning. For the example at hand, we will focus on the idea that the doctor does not aim at causing the pain.

13 Admittedly some may be inclined to say that this example seems unrealistic. It is probably true. But for the sake of the example, let us just assume that stomping on the foot is the only way that Andrew can get Bob to get away from the snake. Or, since Andrew needs to think fast and does not have time to consider other options, we can also assume he justifiably sees it as the best option.

Wesley H. Bronson

responsible for the pain. And as we will see, his being morally responsible results in a counterintuitive conclusion. Some may argue that regardless of the positive gain of saving Bob's life, Andrew should nevertheless be held morally responsible for the pain. Given the situation, Andrew has infinite possibilities, but in terms of his approach to Bob, he merely has two. On the one hand, he can choose to stomp on Bob's foot, intentionally causing pain and thus saving Bob's life. On the other hand, he can perform any other action, all of which will result in Bob's being bitten by the snake and his subsequent death. Given Andrew's ability to either save Bob's life or permit Bob to die, how ought we to look at his choice of action? Admittedly, some may debate the difference between causing some harm and letting some harm merely occur. In this situation, Andrew could actively stomp on Bob's foot, or he could do nothing and permit him to be bitten by the snake. Someone could therefore plausibly argue that if Andrew actively caused the harm, he can be held responsible, while if he permitted Bob to get bitten by the snake, he would not be causing any harm. I do not agree with this claim, nor will I debate the issue over causing a harm to occur versus failing to prevent a harm. For the case at hand, the salient feature of the example is that a harm will occur regardless of what Andrew does. What is up to him, however, is which harm will occur. To save Bob's life, Andrew needs to perform only a small action which will, unfortunately, cause Bob harm. Because the consequence of saving Bob's life so greatly outweighs the minor inconvenience of the pain Bob will incur for that consequence to obtain, it should be obvious that Andrew ought to cause Bob the minor harm and save his life. And it would seem wrong to hold Andrew blameworthy for Bob's pain, for in causing the minor, temporary harm, he saved Bob's life. This, however, is not to say that he is not morally responsible for the pain. He may be morally responsible but nevertheless not be blameworthy, for we commend his saving Bob's life. It may even be that if Andrew ought to cause the harm, then he is positively morally responsible for doing so.

Perhaps then Andrew is doomed to be morally responsible for his choice regardless of what outcome ensues. He can be both causally responsible and morally responsible in this example. It may be asked: regardless of the choices he had to make, did Andrew not aim at stomping on Bob's foot and intentionally causing him pain? It might seem that we are forced to concede that Andrew *is* morally responsible, but he made a good choice nevertheless. If he chooses to stomp on Bob's foot, then he intentionally caused Bob pain. He should therefore be morally responsible

Intentionality, Intending and Moral Responsibility

for the consequence. Similarly, if permitting something bad to happen is as bad as causing it to occur—and we know that Andrew can easily prevent a terrible thing from occurring—we hold Andrew responsible for Bob's death through snake bite even though, as previously mentioned, he is not blameworthy. In our case then, Andrew might be held morally responsible for either consequence that occurs. What is striking here is that we tend to think that if an agent had alternate possibilities and is morally responsible, then he is either praiseworthy or blameworthy. And we have said that Andrew should be commended for causing Bob the minor harm, not criticized for his action. So even though he is morally responsible, he is nevertheless praised for his act. If Andrew is morally responsible, his being praiseworthy is itself a counterintuitive conclusion of the example at hand, possibly negating an idea that we generally take for granted—that if I'm morally responsible for causing pain then I'm not praiseworthy for causing pain.

But we are not committed to the above claim, however, for Andrew is not necessarily morally responsible for causing Bob the minor harm. It may be argued that Andrew is not morally responsible for either action he performs. If we assume that Andrew is a morally upstanding individual and has lived a moral life, yet somehow circumstances put him in this situation, and we recognize that he did not knowingly put himself in this situation, Andrew is an innocent individual. Rather than hold him responsible for either consequence, he ought to be innocent regardless of which choice he makes, for it is merely the circumstances that force him to choose between one of two bad outcomes. Andrew recognizes that it is unfortunate he must make this choice between causing Bob a harm at the least, or observing his death, so he intends for Bob to not experience any pain at all. This, however, is not an option. Given the situation, therefore, we may commend Andrew for saving Bob's life in causing the pain, but we cannot hold him morally responsible for causing Bob the minor harm, even if it was performed intentionally.

If we are right that Andrew intentionally causes Bob's pain and is causally responsible, but is not morally responsible, what other condition would have to hold for Andrew to be morally responsible for the consequence? What other feature of Andrew's mental state would we have to add so that we can unquestionably hold him responsible? It should be apparent that only if Andrew desired for Bob to experience the pain would he be responsible for stomping on his foot. In the situation above, we assume that Andrew does not wish for Bob to

Wesley H. Bronson

experience any pain at all, but since Bob's only chance at survival forces Andrew to intentionally cause him pain, Andrew cannot avoid a harmful consequence. But if Andrew gets some enjoyment out of Bob's pain, then the situation is very different. Andrew's intending to cause Bob pain in this situation is a necessary condition for moral responsibility. What previously vindicated Andrew's action in the seeming dilemma was that although he intentionally brought about a consequence, he did not want to cause the pain in the first place. So if we remove that component of Andrew's mental state and turn him instead into an agent who desires Bob to experience the pain of having his foot stomped on, then we are inclined to hold him morally responsible for Bob's minor harm. Indeed, what vindicates moral responsibility is the lack of intending a specific consequence even if the action is performed intentionally.

Mele and Sverdlik raise an issue that makes us question when we consider an action to have been performed intentionally, and thus force us to look at another type of case for moral responsibility. Referring to a problem first raised by Ronald Butler, they quote the following problem:

If Brown in an ordinary game of dice hopes to throw a six and does so, we do not say that he threw the six intentionally. On the other hand if Brown puts one cartridge into a six chambered revolver, spins the chamber as he aims it at Smith and pulls the trigger hoping to kill Smith, we would say if he succeeded that he had killed Smith intentionally. How can this be so since the probability of the desired result is the same?¹⁴

Butler's problem is not entirely complicated. Brown has a one in six probability that the dice will land six up, and a one in six probability that the gun will shoot a bullet and kill Smith. To make the problems closer to one another, Mele and Sverdlik create their own situations similar to Butler's case. In the first, Brown wants to kill Smith (K-ing). In order for this to happen, however, Brown must throw a dice so that it lands six face-up, causing a chain of events which he knows will result in an explosion that kills Smith. So in throwing the dice, Brown has a one in six probability that he will accomplish his goal and Smith will die. In the second case—which is supposed to more closely resemble the non-moralistic case of simply getting a six in a game of dice—Brown wants to gain membership to a given fraternity (G-ing). Like in the first case,

¹⁴ Ibid., 279.

Intentionality, Intending and Moral Responsibility

Brown can only gain membership by rolling a six, so the chance that he gains membership is one in six. The control Brown has in each case is the same, so is it possible that one action was performed intentionally and the other was not?

Mele and Sverdlik's cases are accurate parallels of Butler's original problem. So Butler might be inclined to say that Brown did not intentionally G but he did intentionally K. Why might he be inclined to think this? Since there is only a limited control Brown has over the possibility of throwing a six, it is difficult to say that he did it intentionally. Yet Butler claims we are inclined to say that Brown killed Smith intentionally. In light of Mele and Sverdlik's examples, however, we see that there is no reason to be committed to this claim. That is, there is no reason to think that one action is intentional and the other is not. We must be either committed to the fact that Brown neither K'd nor G'd intentionally, or that he did intentionally perform both actions. As they point out, there is nothing in either of the cases to make it seem that it is plausible that one was more likely to be performed intentionally than the other. Brown, in each case, merely threw the dice hoping for a certain outcome, but while having equal control over it. In the end, however, they claim that Brown ought to be held responsible for the death of Smith, but that the moral responsibility does not lie in the fact that he performs the action intentionally. In doing so, however, they miss the most salient argument for why we hold Brown responsible. So what do Mele and Sverdlik overlook?

The answer lies in a deeper analysis of intentionality and intending. What they ignore here is *why* we are still inclined to hold Brown responsible, aside from the fact that he is partly causally responsible. As I argued earlier, I maintain that intending is a necessary condition for moral responsibility. What seems most important in the Brown case is that Brown *intended* for the death of Smith to occur. And although he had only minimal control over the outcome, he was as causally responsible for the event as he could be. It is the fact that he intended the outcome that causes him to be responsible for the death. Similarly therefore, Brown is responsible for throwing a six, whether or not he intentionally did it. In killing Smith, Brown intended to throw a six. We again support the thesis, therefore, that under certain situations, intending a consequence is a necessary condition for moral responsibility.

Someone can respond to all of this, of course, by claiming that even if Brown, for example, did not intend to bring about Smith's death,

Wesley H. Bronson

his rolling the die intending for Brown to survive makes him no less blameworthy for Smith's death. Perhaps then Brown is responsible for the action that he does not intend, and I am wrong to claim intending is necessary for moral responsibility. On the one hand, this may seem like a plausible defense. It seems like we have created a situation in which Brown is morally responsible despite not intending for the consequence to obtain. But we must ask a question about an agent who performs an action knowing a certain consequence may occur but performs the action hoping for it not to occur. Namely, why did he perform the action in the first place? If it turns out that Brown had no reason at all to kill Smith and he did not want Smith to die, then there was no reason to roll the dice knowing that Smith might die. If this is the case, then of course we should hold Brown responsible. Performing some action that we know will lead to a consequence simply on a whim—even if we don't intend that consequence—cannot negate moral responsibility. On the other hand, if Brown were to throw the dice intending to throw a five, but had some other motive for throwing the dice, for example the possibility of some greater good, then we would not hold Brown responsible. As mentioned earlier, even if an agent is morally responsible, he need not be blameworthy for the consequence. Andrew might be morally responsible for the pain Bob experiences in having his foot stomped on, but Andrew is not blameworthy for causing the pain. So it is still plausible that even if Brown intends to throw a five, he can still avoid being blameworthy, even if he is causally responsible and intentionally brings about Smith's death.

At this point, I hope it has been shown that under certain circumstances, being causally responsible and intentionally bringing about a consequence is insufficient for holding the agent morally responsible for that consequence. In these situations, what enables us to hold the agent responsible is only when he intends that particular consequence. Let us now look back at the double effect bomber example. Under more careful scrutiny, it is no different than the doctor example. In each situation, the agent is causally responsible for the consequences of his actions. And both the doctor and the bomber intentionally cause some harm. The doctor intentionally causes his patient's pain, and the bomber intentionally kills the innocent civilians. But if I have shown thus far that doing something intentionally and being causally responsible for the consequences are insufficient conditions for moral responsibility, then the bomber is in the same position as was the doctor. And if we claim that the doctor is not morally responsible for the pain his patient incurs,

Intentionality, Intending and Moral Responsibility

then similarly the bomber cannot be held responsible for the death of the innocent civilians. He is causally responsible and intentionally kills the innocent civilians, but he does not intend to cause their death, and as we have shown, when these conditions are satisfied, the agent is not morally responsible. Only when the bomber intends to cause the death of the innocent civilians and then bombs the munitions plant can he be held morally responsible for his actions.

The extent to which this idea can be pushed further, however, is unclear. Let's say that a new poison was found, but it was unknown how long it takes to kill an individual. So for the sake of knowing this information, I administer the poison to someone. I do not want the person to die, but know it is an inevitable outcome of administering the poison. And the only way to know how long the poison takes to kill the man is to perform this experiment. In this situation, I am causally responsible for the death of the man, I intentionally kill him, but I nevertheless do not intend for him to die. If there were a way to find out how long the poison takes to work without killing him I would perform that experiment instead, but unfortunately, there is no other option. While I wish the man could survive, if I am to find out how long the poison takes to work, I must intentionally kill him. At first, it seems that I must be morally responsible for the death of the man. After all, I poisoned him simply to find out how long the poison would take to work. There is no redeeming feature here (for example, that his death will allow me to create an antidote that will save hundreds). Simply put, I kill him to find out some useless information. Can I really be morally innocent here? Admittedly it feels wrong, but if what I have said previously holds true, I am forced to accept that I am not morally responsible for the man's death. While someone may argue that we are inclined to accept the bomber as not morally responsible for the death of the innocents because destroying the munitions plant seems to be an acceptable cost, such a defense would miss the purpose of this paper. Appealing to consequentialist reasoning would precisely avoid the jarring nature of the counterintuitive conclusions about the relationship between intentionality, intending and being morally responsible. The fact that the destruction of the munitions factory is a worthwhile target does not vindicate the bomber in killing the innocents. In fact, a consequentialist could care less about this entire paper's discussion on intention. So appealing to such reasoning defeats the purpose of differentiating doing something intentionally and intending a consequence. We vindicate the bomber on the basis that he

Wesley H. Bronson

did not intend to cause the death of the innocent civilians. The fact that some greater good could be achieved through his action is not germane to the debate. It is simply because although he intentionally kills them, he does not intend to kill them. So while I intentionally kill the man to find out how long the poison takes to work, it is entirely plausible that if I see his death as an unfortunate consequence of the poison experiment and do not intend to bring about his death but know it will occur, I am not morally responsible.

In the end, those who claim that an agent who brings about a known consequence ought to be morally responsible for that action ignore a major feature of the mental state of the agent. For there are times when we do not want to perform an action or do not want to bring about a given consequence, but know it must be done. Under these situations, while we intentionally bring about the known consequence, it seems we can be innocent. But while we may have found a plausible argument for allowing an agent to perform intuitively horrible acts and nevertheless be innocent, it does not seem that we must allow the agent to remain completely blameless. After all, if I kill a man for the sake of finding out how long a poison will take to work, regardless of whether you hold me morally responsible or say there is a redeeming feature—that I did not intend the consequence—I did nevertheless kill the individual. Some may be inclined to say that the position is implausible. We have a universal intuition that the individual who kills a man as an unintended consequence of some trivial act he did intend is definitely morally responsible. There are two responses to such an objection. First, the intuition about moral responsibility lacks a certain normative force. It is not so clear why these intuitive gut responses to the conclusions of the paper provide a justifiable objection. Second, the intuitions themselves seem linked to consequentialist reasoning. The detractor of the view presented here will respond by noting how the cost of the death cannot possibly outweigh the minor gain of the experiment. If we are to support a degree of consequentialist reasoning in our deliberations, then the objection would be strong. But if we do not include such reasoning, then it is the burden of the detractor to explain what it is specifically that makes the view implausible.

Intentionality, Intending and Moral Responsibility

References

- Bentham, Jeremy. *An Introduction to the Principles of Morals and Legislation*. Ed. J Burns and H. Hart. London: Methuen, 1970.
- Bratman, Michael. *Intention, Plans, and Practical Reason*. Cambridge: Harvard University Press, 1987.
- Frankfurt, Harry G. "Alternative Possibilities and Moral Responsibility." *The Journal of Philosophy*. Vol. 66, No. 23, (Dec 4, 1969), pp. 829-839
- Ginet, Carl. "The Epistemic Requirements for Moral Responsibility." *Philosophical Perspectives 14: Action and Freedom*, 2000. Ed. James E. Tomberlin. Boston: Blackwell Publishers, 2000. 267-277.
- Mele, Alfred and Steven Sverdlik. "Intention, Intentional Action, and Moral Responsibility." *Philosophical Studies* Vol. 82, (1996) pp. 265-287.