

Desire, Foresight, Intentions, and Intentional Actions: Probing Folk Intuitions

Thomas Nadelhoffer

Florida State University

Abstract:

A number of philosophers working under the rubric of “experimental philosophy” have recently begun focusing on analyzing the concepts of ordinary language and investigating the intuitions of laypersons in an empirically informed way.¹ In a series of papers these philosophers—who often work in collaboration with psychologists—have presented the results of empirical studies aimed at probing folk intuitions in areas as diverse as ethics, epistemology, free will, and the philosophy of action. In this paper, I contribute to this research program by discussing the results of some new experiments that further probe folk intuitions about the relationship between desire, foresight, intent, intentional action, and moral considerations.

Keywords:

folk intuitions, moral judgments, foresight, intention, intentional action

Introduction

Philosophers have long been interested in the relationship between desire, foresight, intentions, and intentional actions—and for good reason. After all, on a number of occasions, determining how much praise or blame an agent deserves may depend on our having first understood the proper connection between this closely related cluster of concepts. For instance, should we hold an agent who knowingly but reluctantly brings about an undesired side effect *y* just as morally or legally responsible as someone who actually aims to bring *y* about? Adequately answering this question will depend in part on how one analyzes the connection between the concepts of desire, foresight, intention,

¹ My goal in this paper is not to suggest that one of these approaches is superior to the other. Indeed, on my view, both approaches have their respective merits and drawbacks. For present purposes, I will simply assume that understanding the relationship between the folk concepts of desire, foresight, intention, and intentional action is an important philosophical goal in its own right. If nothing else, it seems clear that our ordinary concepts should be the starting point of conceptual analysis—even if it remains open that once we learn more about these concepts, we may very well find it necessary to either revise or replace them.

and intentional action—which will in turn depend in part on what one takes to be the proper subject matter of conceptual analysis as it relates to the philosophy of action.

Some philosophers are primarily interested in developing explanatory models of the etiology of intentional actions—and hence they focus on providing analyses of technical philosophical concepts and investigating the psychological events and processes that play a role in the production of intentions and intentional actions. And while there is certainly much to be said for this approach to conceptual analysis, a number of philosophers working under the rubric of “experimental philosophy” have begun focusing instead on analyzing the concepts of ordinary language and investigating the intuitions of laypersons.² In a series of recent papers these philosophers—who often work in collaboration with psychologists—have presented the results of empirical studies aimed at proving folk intuitions in areas as diverse as ethics (Doris and Stich forthcoming; Nichols 2004), epistemology (Nichols, Weinberg, and Stich 2002), free will (Nahmias et al 2004; Nichols 2004), and the philosophy of action (Knobe 2003a; 2003b; 2004; Malle and Knobe 1997; Nadelhoffer 2004; forthcoming).

One of the more important and interesting findings to come from the empirical research related to the philosophy of action is the fact that moral considerations often influence folk ascriptions of intentional action. More specifically, it appears that people are more likely to say that an agent intentionally performed a morally bad action than they are to say that an agent intentionally performed a structurally similar morally good

² My goal in this paper is not to suggest that one of these approaches is superior to the other. Indeed, on my view, both approaches have their respective merits and drawbacks. For present purposes, I will simply assume that understanding the relationship between the folk concepts of desire, foresight, intention, and intentional action is an important philosophical goal in its own right. If nothing else, it seems clear that our ordinary concepts should be the starting point of conceptual analysis—even if it remains open that once we learn more about these concepts, we may very well find it necessary to either revise or replace them.

or neutral action (Knobe 2003a; 2003b; 2004; Nadelhoffer 2004; forthcoming). This finding goes right to the heart of a long-standing debate in the philosophy of action concerning whether moral considerations do – and whether they should – affect our application of the concept of intentional action (see, e.g., Bratman 1987; Butler 1978; Duff 1982; 1990; Harman 1976; Katz 1987).

The gathering data about moral judgments and ascriptions of intentional action also shed light on the task of giving a precise psychological characterization of the nature of the boundaries between folk psychology and folk morality. In this paper, I try to contribute to the philosophical debate and the psychological task by discussing the results of some new experiments that further probe folk intuitions about the relationship between desire, foresight, intent, intentional action, and moral considerations.

Throughout this paper I operate under the assumption that one way to test whether particular analyses of intention or intentional action are in line with ordinary usage is to treat them as predictions about how laypersons would respond to particular cases. As we will see in §§2-4, I have employed this method in my efforts to probe folk intuitions concerning some famous examples from the action theory literature. My goal in running these studies was not to determine whether particular analyses are *true*—rather I was merely trying to determine whether these analyses cohere with our ordinary beliefs, judgments, and intuitions about the relationship between desire, foresight, intentions, intentional actions and moral responsibility.³ Moreover, I hope to provide evidence

³ By my lights, until this preliminary step is taken we will not be in a proper position to determine the true merits of any of the analyses of intention and intentional action that I examine in this paper. After all, while philosophical analyses need not be entirely beholden to folk concepts and ordinary intuitions, they should not be entirely unanchored by them either—especially in areas of philosophy such as action theory that deal with concepts that are central to our everyday moral and legal practices.

concerning the aforementioned issue of the boundaries of folk psychology and folk morality.

1. Setting the Stage

Owing to the important role that ascriptions of intentional action play in our daily lives, it is not surprising that the concept of intentional action often takes center stage in discussions among action theorists. But what does it mean to act intentionally and what is the relationship between intending to x and intentionally x -ing? According to the analysis of intentional action that Michael Bratman has dubbed the “Simple View (SV),” in order for an agent to intentionally x , she must intend to x (Bratman 1984: 377). On this view, in order for my going to the park to count as an intentional action, at the time that I go to the park my mental states must be such that going to the park is among the things that I intend to do (see, e.g., Adams 1986; McCann 1986; 1991).⁴

Proponents of this view defend it on a number of grounds. First and foremost, the SV purportedly captures our pre-theoretical intuitions and is in line with our ordinary usage of the concepts of intention and intentional action (McCann 1998: 210). After all, in ordinary contexts it would admittedly sound strange for me to say that I dialed my friend’s phone number intentionally even though I did not intend to do so. Second, to the extent that the SV is the seemingly uncontroversial claim that *intending* to x is necessary for *intentionally* x -ing, the view has the virtue of being—well, simple or “uncluttered” (Adams 1986: 284). Finally, the SV “gives us reason to believe that our intentions

⁴ While proponents of the SV do claim that an intention to x is necessary for x -ing intentionally, they do not make the further claim that an agent who intends to x and actually does x , does x intentionally. On their view, the intention must cause the action “in the right way” for the action to be intentional (Adams 1986: 284).

causally guide our actions in virtue of their content” (Adams 1986: 284)—thereby supporting our ordinary view of ourselves whereby the contents of our intentions to *x* play an important role in our intentionally *x-ing*.

However, not all philosophers are proponents of the SV. Gilbert Harman, for instance, has suggested that, “it is a mistake to suppose that whenever someone does something intentionally, he intends to do it” (1976: 433). To support this claim, Harman points to cases involving foreseen yet unintended and undesired side effects. By his lights, for instance, a sniper may intentionally alert his enemies when he shoots his primary target, even though he neither desires nor intends to alert them (1976: 433). On this view, we can properly say that the sniper intentionally alerted the enemies even though he did not intend to because “in firing his gun, the sniper knowingly alerts the enemy to his presence. He does this intentionally, thinking that the gain is worth the possible cost. But he certainly does not intend to alert the enemy to his presence” (1976: 433). Harman does not want to say that the soldier intended to alert the enemies because doing so was something he neither tried nor wanted to do. Indeed, the soldier had both a desire and a reason *not* to alert them.

According to Harman’s account of the relationship between intention and intentional action, “one can do something intentionally even though one does not intend to do it, if one does it in the face of what ought to be a reason not to do it and, either one tries to do it, or one does it as a foreseen consequence of something else that one intends to do” (1976: 434). To highlight this aspect of his analysis, Harman claims that in addition to knowingly alerting his enemies to his presence, the sniper also knowingly heats the barrel of his gun by firing it. However, according to Harman, this latter side

effect, unlike the former, is not intentionally brought about because the sniper does not have a reason not to heat the barrel of his gun in the way that he has a reason not to alert his enemies.

On this “reasons against” view, if an agent (a) foresees that by doing *x* she will bring about *y*, (b) she has a reason not to bring *y* about, and (c) she does *x* intentionally, then she brings about *y* intentionally as well, even if she neither wanted, aimed, nor intended to bring *y* about. Indeed, it is precisely this alleged asymmetry between what is required for intending to *x* and what is sufficient for intentionally *x-ing* that allows Harman to use the case of the sniper as a counter-example to the SV. After all, the sniper intentionally brought something about that he did not want to bring about—namely, alerting his enemies to his presence. And given that he did not want to alert his enemies, he did not intend to do. If Harman is correct, the sniper intentionally alerted his enemies without intending to do so—thereby falsifying the SV.

Proponents of the SV cannot simply deny that Harman’s sniper intentionally alerted the enemies to his presence on the grounds that the sniper did not intend to do so without begging the very question that is it at issue. However, there is another way for them to respond—namely, they could suggest contra Harman that the sniper intended to alert the enemies to his presence after all on the grounds that doing *x* with the foresight that *x* will produce *y* may be sufficient for both intending to *y* and intentionally bringing *y* about so long as other salient conditions are met (e.g., doing *x* was not the result of pure luck and bringing *y* about was not the result of causal deviance). If this response to Harman works, then the SV will be insulated from his alleged counter-example.

Jeremy Bentham was one of the first philosophers to argue in detail that if an agent intentionally does x knowing that x will bring about y , then she both intends to bring y about and brings y about intentionally even if y is a state of affairs that the agent does not want to bring about. Moreover, Bentham suggests that in order for either x or the consequences of x to count as intentional, the agent must have either intended to x or intended to bring about the consequences of x (1789: 84). In at least this one respect, Bentham's analysis of the relationship between intention and intentional action is in line with the SV.

For present purposes, the most important part of Bentham's theory of action is the distinction between "direct" and "oblique" intentions—a distinction that is drawn in the following way: When the prospect of bringing about a set of consequences "constituted one of the links in the chain of causes by which the person was determined to do the act," then these consequences are brought about directly or "lineally" (1789: 86). On this view, if bringing y about was part of an agent's motivation for doing x , both x and y were directly intended and hence y was brought about intentionally. On the other hand, when an agent merely foresaw but did not desire that y would occur as a result of doing x —i.e., when the prospect of bringing y about did not constitute one of the reasons for which the agent x -ed—bringing about y is something the agent intended—albeit obliquely or "collaterally" rather than directly (1789: 86). And since in this latter case the agent obliquely intended to y , the agent still brought y about intentionally even though she neither wanted nor tried to do so.

In arguing for the correctness of the distinction between direct and oblique intention, Bentham introduces his famous stag-hunter example (1789: 87). We are asked

to imagine a stag-hunter who spots a stag in the woods. But before the hunter shoots an arrow at the stag, he notices that the King is in close proximity. And while the hunter does not want to shoot the King, he realizes that if he shoots the stag, he will likely kill both the stag and the King. This realization notwithstanding, the hunter shoots and kills the stag--thereby shooting and killing the King in the process. According to Bentham, while the hunter directly or lineally intended to shoot and kill the stag given that doing so constituted part of his reason for acting, he merely obliquely intended to shoot and kill the King given that he foresaw that he would shoot the King but did not want or aim to, do so. In any event, insofar as the death of the stag and the death of the King were both intended—albeit to varying degrees—they were both brought about intentionally (1789: 87). If Bentham’s analysis is correct, then Harman’s sniper intended to alert his enemies after all. And as we saw earlier, this would thereby undermine the usefulness of these examples for refuting the SV. Of course, not everyone accepts Bentham’s analysis.

One philosopher who disagrees with Bentham on this score is Elizabeth Anscombe. In her groundbreaking work *Intention* (1957), Anscombe claims that doing x with the foresight that x will bring about y is *not* sufficient for either intending to y or intentionally y -ing. On her view, wanting to bring y about is necessary for intending to y and for bringing about y intentionally. The example that Anscombe uses to defend this claim has proved to be no less controversial than Bentham’s own example. We are asked to imagine a man whose job it is to pump water into a cistern, thereby replenishing the supply of drinking water in a nearby house (1957: 37). Unfortunately for the inhabitants of the house, the water that the man is pumping into the cistern has been contaminated with a lethal poison whose effects are unnoticeable until those who ingest it can no longer

be cured. Even though the man pumping the water had nothing to do with poisoning the water, he knows that the water has been poisoned. Nevertheless, the man pumps the water into the cistern knowing that it will poison and kill the inhabitants. But, he neither wants nor aims to kill them; he just wants to do his job and get paid. He views the death of the inhabitants as an unfortunate by-product of his pumping water into the cistern.

The question this example is supposed to highlight is the following: Can we properly say that the water pumper did not intend to kill the inhabitants even though he foresaw that pumping the water into the cistern would inevitably have this effect? Anscombe answers this question in the following way:

A further difficulty arises from the fact that the man's intention might not be to poison them but only to earn his pay. That is to say, if he is being improbably confidential and is asked "Why did you replenish the house water-supply with poisoned water?," his reply is not "To polish them off," but "I didn't care about that, I wanted my pay and just did my usual job." In that case, although he knows concerning an intentional act of his...that it is also an act of replenishing the house water supply with poisoned water, it would be incorrect, by our criteria, to say that his act of replenishing the house water supply with poisoned water was intentional. And I do not doubt the correctness of this conclusion; it seems to show that our criteria are rather good. (1957: 41-42)

Not surprisingly, some philosophers have taken Anscombe's intuitions about the water pumper case to show that her criteria for intentions and intentional actions are *actually rather bad*.

Arthur Miller, for example, suggests that Anscombe's claim that "desirability characterizations" (1957: 75) are necessary for both intending to x and intentionally x -ing is "unable to withstand analysis" (Miller 1980: 337). By his lights, "the pumper did intend to poison the inhabitants as well as replenish the water supply; or alternatively, to establish that poisoning the inhabitants was as much an intentional action on the part of the pumper as was replenishing the water supply" (1980: 337). In arguing against

Anscombe's view, however, Miller does not deny that there is a close connection between intending to x (or intentionally x -ing) and wanting to x —he merely denies that the former entails the latter. In his efforts to show that wanting is not necessary for either intention or intentional action, Miller discusses what he takes to be the implications of Anscombe's view—implications that he finds “clearly unacceptable” (1980: 339).

According to Miller, one of the most important factors that we take into consideration when making moral judgments about an agent's actions (or the consequences of those actions) is the question of whether or not the agent acted intentionally. As he says, “we behave differently toward an agent whose act is intentional than we do toward an agent whose act is not intentional” (1980: 340). Thus, if A killed B intentionally whereas C killed D but not intentionally, then all other things being equal, we will judge A more harshly than C. In other words, if an agent either unintentionally or non-intentionally x -ed, this ordinarily serves either as an excusing condition or, minimally, as a mitigating circumstance. Conversely, we normally hold an agent fully responsible for x -ing if she either intended to x or x -ed intentionally. And while Miller allows that intending to x often involves desirability characterizations, he claims that these characterizations are by no means necessary. On his view, “if Y is not only the foreseeable result of doing X in a given set of circumstances, but foreseen by the agent himself, then if doing X is to count as an intentional action on his part, then so is doing Y” (1980: 341).

Miller suggests that if his analysis of the relationship between foresight and intentional action were not correct, we would not be able to hold Anscombe's water pumper *fully* responsible for killing the inhabitants given that lack of intention is either an

excusing condition or a mitigating circumstance—all other things being equal. As he says, “if we know that A was willing to do Y and he then went on to do it, we will certainly hold him responsible and...blame him for Y-ing; and *ceteris paribus* this we would not do if Y were not intentional” (1980: 342). After all, by Miller’s lights, if intending entails wanting, and Anscombe’s water pumper neither wanted to nor cared about poisoning the inhabitants, then he did not intend to poison them. And to the extent that Miller thinks this would “incline us to view the episode more leniently than we might otherwise have done”—a consequence that he finds “patently absurd” given the circumstances—he concludes that Anscombe’s view that intending entails wanting must be incorrect (1980: 341).

For present purposes, I am not going to try and settle any of the issues I have discussed in this section. Instead, I want to turn our attention to the results of some attempts to test whether any of the analyses of intention and intentional action that we have covered are in line with the ordinary intuitions of laypersons. As I suggested earlier, one way to test whether particular analyses of intention or intentional action cohere with ordinary usage is to treat them as predictions about how laypersons would respond to particular cases. So, I conducted some simple surveys in order to probe folk intuitions about the sorts of famous cases we have just discussed from the action theory literature. Now it is time to examine the results.

2. Bentham’s Stag-Hunter Revisited

As we saw earlier, Bentham introduced his famous stag-hunter case in an effort to illustrate his distinction between oblique and direct intentions. On his view, the stag

hunter intended to kill the king—albeit obliquely rather than directly—given that the king’s death was a foreseen, but undesired, side effect. If this claim is taken to be one about our ordinary usage of the concepts of intention and intentional action, then we can treat it as a prediction about how laypersons would respond to Bentham’s example.⁵ To test this prediction, I ran a series of surveys. Participants were 160 undergraduates—each of whom received one of the following vignettes:

Case 1 (C1):

Imagine that there is a man out in the woods who is participating in a hunting competition. After spending hours waiting for a deer to cross his path, the hunter suddenly sees the largest deer he has ever seen. If he can only kill this deer, he will surely win the competition. So, the hunter gets the deer in his sights—but at the last second, he notices that there is a group of bird-watchers just on the other side of the deer. The hunter realizes that if he shoots the deer, the bullet will *definitely* hit one of the bird-watchers as well. But he does not care at all about the bird watchers—he just wants to win the competition. So, he shoots and kills the deer. And as expected, the bullet ends up hitting one of the bird-watchers as well.

Case 2 (C2):

Imagine that there is a man out in the woods who is participating in a hunting competition. After spending hours waiting for a deer to cross his path, the hunter suddenly sees the largest deer he has ever seen. If he can only kill this deer, he will surely win the competition. So, the hunter gets the deer in his sights—but at the last second, he notices that there is a group of bird-watchers just on the other side of the deer. The hunter realizes that if he shoots the deer, the bullet could *possibly* hit one of the bird-watchers as well. But he does not care at all about the bird watchers—he just wants to win the competition. So, he shoots and kills the deer. And as expected, the bullet ends up hitting one of the bird-watchers as well.

Case 3 (C3):

Imagine that there is a man out in the woods who is participating in a hunting competition. After spending hours waiting for a deer to cross his path, the hunter suddenly sees the largest deer he has ever seen. If he can only kill this deer, he will surely win the competition. So, the hunter gets the deer in his sights—but at the last second, he notices that there is a group of bird-watchers just on the other side of the deer. The hunter realizes that if he shoots the deer, there is a *very small chance* that the bullet will hit one of the bird-watchers as well. But he does not care at all about the bird watchers—he just wants to win the competition. So,

⁵ This is not to suggest that Bentham either believed or claimed that his view settled with our pre-theoretical intuitions—rather, I am simply trying to determine whether or not it actually does.

he shoots and kills the deer. Unfortunately, the bullet ends up hitting one of the bird-watchers as well.

Case 4 (C4):

Imagine that there is a man out in the woods who is participating in a hunting competition. After spending hours waiting for a deer to cross his path, the hunter suddenly sees the largest deer he has ever seen. If he can only kill this deer, he will surely win the competition. So, the hunter gets the deer in his sights and pulls the trigger—thereby killing the deer. Unfortunately, the bullet exited the deer’s body and struck a hunter who was hiding nearby.

Participants were then asked the following six questions:

- 1) Did the man intentionally shoot the bird-watcher [hunter]?
- 2) Did the man want to shoot the bird-watcher [hunter]?
- 3) Did the man intend to shoot the bird-watcher [hunter]?
- 4) Did the man purposely shoot the bird-watcher [hunter]?
- 5) How much blame does the hunter deserve for shooting the bird-watcher [hunter] (On a scale from 0 to 6—0 being no blame and 6 being a lot of blame)?

The results were as follows:

ANSWERS:	C1:	C2:	C3:	C4:
1) Intentionally	68%	30%	35%	0%
2) Wanted	13%	0%	0%	0%
3) Intended	45%	5%	5%	0%
4) Purposely	60%	25%	15%	0%
5) Blame rating	5.84	5.5	4.5	2

One noteworthy feature of these results is that they support the view that an agent can intentionally bring a side effect about even though she did not intend to do so. After all, there is a statistically significant difference between participants’ responses to questions (1) and (3) in C1-C3 (χ^2 (1, N=120) = 19.3, $p < 0.001$). Given that more participants judged that the hunter intentionally shot the birdwatcher than judged that the hunter intended to shoot her, it appears that we cannot simply read off an agent’s intentions from her intentional actions in the way Bentham assumes. And to the extent that people think

that we can intentionally x without intending to x , this speaks against the SV as an analysis of the ordinary concepts of intention and intentional action as well.⁶

A second noteworthy feature of these results is the surprising fact that more participants in C1 judged that the hunter intended to shoot the birdwatcher [or hunter] than judged that the hunter wanted to bring this side effect about. This difference between participants' ascriptions of intention and their ascriptions of desire is statistically significant (χ^2 (1, N=40) = 10.3, $p < 0.01$).⁷ The fact that this asymmetry exists is important given that a number of philosophers (e.g., Audi 1973; 1986; Austin 1873; Davis 1984; 1997) ascribe to a “desire-belief” model of intention whereby in order for an agent to intend to perform some action x , she must satisfy some minimal cognitive and conative conditions. On this view, intending to x requires that, (a) the agent either knows or believes that she will x (or likely x), and (b) the agent desires, wants, or prefers to x . John Austin, for example, defines “a present intention to perform a future action,” as, “a present *desire* of an object (either as an end or as a means), coupled with a present belief that we shall do acts hereafter for the purpose of attaining the object” (1873: 451). The results of C1 (and to a lesser extent C2-C4) suggest that people may not appeal to the “desire belief” model when judging particular cases.

⁶ The debate about whether the SV is a correct analysis of the folk concept of intentional action has received a lot of attention in the action theory literature recently. Knobe was the first philosopher to attempt to test whether the SV settled with ordinary intuitions (Knobe 2004). Adams and Steadman then offered some objections to Knobe's experiments and suggested a better way of getting at the relevant intuitions (Adams and Steadman 2004b)—advice that Hugh McCann subsequently incorporated into studies of his own (McCann forthcoming). To his surprise, the results once again suggest that the SV does not settle with ordinary usage. I have a manuscript under review entitled “On Saving the Simple View” where I provide an overview of the relevant issues and studies and where I respond to the various attempts on behalf of proponents of the SV to rescue it from the empirical data.

⁷ When we compare participants' responses from C1-C4 to questions 2 and 3, we get the following result: (χ^2 (1, N=160) = 11.68, $p < 0.001$).

Another interesting—albeit not surprising—feature of these results is that participants’ ascriptions of intentional action and their attributions of blame are sensitive to how confident the hunter was that he would end up shooting one of the bird-watchers. Participants who read that the hunter knew that in shooting the deer he would *definitely* shoot one of the bird-watchers, were roughly twice as likely to judge that he shot the bird-watcher intentionally than those who read that he knew that it was either *possible* or that there was a *very small chance* that he would do so—a highly statistically significant difference ($\chi^2(1, N=120) = 28.51, p < 0.001$).

Moreover, when we compare the way participants in C4—where there was no foresight on the part of the hunter—responded with the way that participants in C1-C3 responded, we find that when an agent foresees a side effect—whether it is perceived to be a definite (C1) or merely a possible (C2 and C3) consequence of one’s actions—the more likely people are to say that it was intended ($\chi^2(1, N=160) = 3.99, p < 0.05$) and/or brought about intentionally ($\chi^2(1, N=160) = 6.47, p < 0.025$). And when we compare the results of the two opposite extremes—namely, C1 where the hunter knew he would definitely shoot a birdwatcher and C4 where the hunter had no idea that he would shoot another hunter—the asymmetry of participants’ intuitions concerning whether the hunter intentionally brought about the respective side effects is further magnified ($\chi^2(1, N=80) = 40.78, p < 0.001$).

Indeed, none of the participants in C4 who read that the hunter did not foresee that in shooting the deer he would be shooting one of the bird-watchers judged that the hunter either intended or intentionally brought about this side effect. And given that in all four cases, the hunter brought about the very same side effect these results suggest that the

badness of the side effect cannot *fully* explain the participants' intuitions—even if moral considerations surely play some role. It seems that the only other variable that sheds light on the asymmetry in participants' responses is the varying degree of the hunter's foresight. Hence, it appears that Bentham was onto something after all—even if his view fails to capture some of our ordinary intuitions. People are obviously reluctant to judge that an agent intentionally brought about an *unforeseen* side effect even if they are not as quick to judge that an agent intentionally brought about a *foreseen* side effect as Bentham may have assumed. In any event, the results of C1-C4 reveal that folk ascriptions of intentional action are at least sometimes sensitive to the perceived likelihood of the side effect. But are moral considerations also playing a part in shaping participants' judgments concerning the intentionality of the consequences of the hunter's actions?

In order to see whether participants' intuitions would change in a similar case that does not involve a morally bad side effect, I ran another experiment. This time, participants were 40 undergraduates—each of whom received the following vignette:

Case 5 (C5):

Imagine that there is a man out in the woods who is participating in a hunting competition. After spending hours waiting for a deer to cross his path, the hunter suddenly sees the largest deer he has ever seen. If he can only kill this deer, he will surely win the competition. So, the hunter gets the deer in his sights—but at the last second, he notices that there is a beautiful eagle perched in a tree nearby. The hunter realizes that if he shoots the deer, the sound of the gunfire will *definitely* cause the eagle to fly away. But he does not care at all about the eagle—he just wants to win the competition. So, he shoots and kills the deer. And as expected, the sound of the gunfire causes the eagle to fly away.

They were then asked the following four questions:

- 1) Did the man intentionally cause the eagle to fly away?
- 2) Did the man want to cause the eagle to fly away?
- 3) Did the man intend to cause the eagle to fly away?
- 4) Did the man purposely cause the eagle to fly away?

And the results were as follows:

ANSWERS:	C5:
1) Intentionally	35%
2) Wanted	10%
3) Intended	5%
4) Purposely	40%

By comparing these results with those of C1, we see that even though in both cases the hunter knew that he would *definitely* bring about the respective side effects, when the side effect was morally negative—i.e., shooting a birdwatcher—nearly twice as many participants (68%) judged that he intentionally brought it about than when the side effect was non-moral (35%)—a highly statistically significant difference ($\chi^2(1, N=80) = 8.46, p < 0.001$). This confirms that moral considerations are at least partially responsible for the results of C1, C2, and C3—even if foresight played a role as well. Moreover, the results of C5 provide us with additional evidence that participants are more likely to judge that the hunter intentionally brought about a side effect than they are to judge that he intended to do so ($\chi^2(1, N=40) = 11.25, p < 0.001$). And since this is precisely the type of response that a proponent of Harman’s analysis of intention and intentional action would predict, we should turn our attention to his view.

3. Harman’s Sniper Revisited

As we saw earlier, according to Harman’s analysis of intention and intentional action, if an agent (a) foresees that by doing *x* she will bring about *y*, (b) she has a reason not to bring *y* about, and (c) she does *x* intentionally, then she brings about *y* intentionally as well, even if she neither wanted, aimed, nor intended to bring *y* about. If we follow the previous protocol and treat this view as a prediction about how laypersons would respond

to particular cases involving agents who satisfy these conditions, then the results of the aforementioned stag-hunter studies lend support to Harman's view. Nevertheless, to test it more directly, I ran another series of surveys based on his aforementioned sniper scenario. This time, participants were 120 undergraduates—each of whom received one of the following three vignettes:

Case 6 (C6):

A sniper has been ordered to kill an enemy commander. So, after getting himself into position, he finally has the enemy commander in his sights. Before he pulls the trigger, however, the sniper realizes that the gunfire will *definitely* alert the other enemy soldiers to his presence. But the sniper doesn't care at all about that—he just wants to shoot his target. So, he pulls the trigger—thereby shooting and killing the commander. And, as he expected, the enemy soldiers are alerted to his presence.

Case 7 (C7):

A sniper has been ordered to kill an enemy commander. So, after getting himself into position, he finally has the enemy commander in his sights. Before he pulls the trigger, however, the sniper realizes that the gunfire will *probably* alert the other enemy soldiers to his presence. But the sniper doesn't care at all about that—he just wants to shoot his target. So, he pulls the trigger—thereby shooting and killing the commander. Unfortunately, the enemy soldiers are alerted to his presence.

Case 8 (C8):

A sniper has been ordered to kill an enemy commander. So, after getting himself into position, he finally has the enemy commander in his sights. Before he pulls the trigger, however, the sniper realizes that there is a *very small chance* that the gunfire will alert the other enemy soldiers to his presence. But the sniper doesn't care at all about that—he just wants to shoot his target. So, he pulls the trigger—thereby shooting and killing the commander. Unfortunately, the enemy soldiers are alerted to his presence.

Then each of the participants were then asked the following four questions:

- 1) Did the sniper intentionally alert the enemies to his presence?
- 2) Did the sniper want to alert the enemies to his presence?
- 3) Did the sniper intend to alert the enemies to his presence?
- 4) Did the sniper purposely alert the enemies to his presence?

The results were as follows:

ANSWERS:	C6:	C7:	C8:
1) Intentionally	55%	40%	20%
2) Wanted	10%	5%	0%
3) Intended	10%	15%	0%
4) Purposely	55%	30%	30%

When we compare the results of C6—where the sniper knew that he would definitely alert his enemies—with the results of C7 and C8—where he foresaw the possibility that he might alert them—we once again find that participants’ judgments concerning whether the sniper intentionally alerted the enemies to his presence were sensitive to the degree of the sniper’s foresight ($\chi^2(1, N=120) = 7.05, p < 0.01$). The less certain the sniper was that he would alert the enemies to his presence, the less likely participants were to judge that he either intended to alert them or that he alerted them intentionally. And, once again, when we compare the way participants in C6-C8 responded to the first and third questions, we have additional evidence that people are often more likely to say that a side effect was brought about intentionally than they are to say that the agent intended to bring it about ($\chi^2(1, N=120) = 46.61, p < 0.001$).

On Harman’s view, the asymmetry between judgments of intention and judgments of intentionality is to be explained on two grounds. First, because the sniper did not *want* to alert his enemies, he did not intend to do so—which would explain why relatively few participants judged that the sniper intended to alert his enemies in any of the cases.⁸ Second, because the sniper shot the enemy commander even though doing so would bring about a foreseen but *undesired* side effect, he intentionally brought this side effect about. Keep in mind that by Harman’s lights in order for an agent to bring about a

⁸ It is worth pointing out that in C6-C8—unlike in C1-C4—there is not a statistically significant difference between the participants’ judgments concerning whether the agent intended and wanted to bring the side effect about.

side effect intentionally, the side effect must be something that the agent would rather avoid, all other things being equal.

To test whether Harman’s “reasons against” view explains the asymmetry in the results of C6-C8, I ran two more experiments. This time, participants were 80 undergraduates—each of whom received one of the following two vignettes:

Case 9 (C9):

A sniper has been ordered to kill an enemy commander. So, after getting himself into position, he finally has the enemy commander in his sights. Before he pulls the trigger, however, the sniper realizes that the gunfire will *definitely* cause the barrel of his gun to get hot. But the sniper doesn’t care at all whether the barrel of the gun is hot, he doesn’t have to touch it anyway. So, he pulls the trigger—thereby shooting and killing the commander. And, as the sniper expected, firing the gun caused the barrel to heat up.⁹

Case 10 (C10):

A sniper has been ordered to kill an enemy commander. So, after getting himself into position, he finally has the enemy commander in his sights. Before he pulls the trigger, however, the sniper realizes that when the bullet travels through the air it will *definitely* disturb the air molecules around it. But the sniper doesn’t care at all about disturbing air molecules, he just wants to shoot his target. So, he pulls the trigger. As the sniper expected, as the bullet travels from his gun into the enemy commander, it disturbed some air molecules.

They were then asked the following four questions:

- 1) Did the sniper intentionally heat the barrel of his gun [disturb some air molecules]?
- 2) Did the sniper want to heat the barrel of his gun [disturb some air molecules]?
- 3) Did the sniper intend to heat the barrel of his gun [disturb some air molecules]?
- 4) Did the sniper purposely heat the barrel of his gun [disturb some air molecules]?

The results were as follows:

ANSWERS:	C9:	C10:
1) Intentionally	68%	45%
2) Wanted	5%	5%

⁹ This scenario is also borrowed from Harman (1976). And as we saw earlier, Harman suggests that we should not say that the sniper intentionally heated up the barrel of his gun given that he does not have a reason *not* to do so in the way that he has a reason not to alert the enemies to his presence.

3) Intended	28%	18%
4) Purposely	60%	35%

When we combine the answers to the first and third questions from C9 and C10, we find that 57% of the participants judged that the sniper intentionally brought the respective side effects about, whereas only 23% judged that he intended to do so—a highly statistically significant difference ($\chi^2(1, N=80) = 19.1, p < 0.001$). In this respect, the results from these two studies are entirely consistent with the results of C1 – C8 when it comes to the falsity of the SV as an analysis of the folk concept of intentional action. However, because the sniper in C9 and C10 does not have any reason to avoid heating up the barrel of his gun or to avoid moving air molecules, we cannot appeal to Harman’s “reasons against” view to explain the asymmetry between participants’ answers to the first and third questions.¹⁰

If this is correct, then it appears that (a) folk ascriptions of intention are more sensitive to conative considerations than their ascriptions of intentional action, (b) both their ascriptions of intention and intentional action are sensitive to moral considerations, and (c) an agent’s degree of foresight plays an important role in people’s judgments concerning whether she intended or intentionally brought about a side effect. Nevertheless, in order for an agent to intentionally bring about a foreseen side effect, she need not have a reason not to do so—at least as far as the majority of the participants in these surveys are concerned.¹¹

¹⁰ It is worth pointing out that there is a statistically significant difference between intentionality judgments in C9 and C10 ($\chi^2(1, N=80) = 4.11, p < 0.05$). I am unsure what explains this difference.

¹¹ It is curious indeed that in these sniper studies—like C1-C4 but unlike C6-C8—we find a statistically significant difference between participants’ judgments concerning whether the sniper intended to bring about the respective side effects and whether he wanted to do so ($\chi^2(1, N=80) = 10.32, p < 0.01$). In order

4. Anscombe's Water Pumper Revisited

The results from both the hunter and sniper studies give us prima facie reason to suspect that Anscombe's view is not in line with ordinary usage either. After all, if we adopt her view, then because neither the hunter nor the sniper wanted to bring about their respective side effects, they neither intended nor intentionally did so. Yet, as we have already seen in the Bentham and Harman studies, laypersons seem to place a higher value on foresight when it comes to judging whether an action was intended or intentional than Anscombe does—especially in cases where the outcome is morally bad. And while I think that the data I have already presented are enough to show that her analysis is inconsistent with the majority of folk intuitions, I ran one more study in order to test her view directly rather than merely indirectly. This time participants were 100 undergraduates—each of whom received one of the following vignettes:

Case 11 (C11):

There is a man who gets paid for pumping water into a cistern thereby replenishing the supply of drinking water in a nearby house. Unfortunately for the inhabitants of the house, the water that the man is pumping into the cistern today has been systematically contaminated with a lethal poison whose effects are unnoticeable until they can no longer be cured. Even though the man pumping the water had nothing to do with poisoning the water, he knows that the water has been poisoned. Nevertheless, the man pumps the water into the cistern knowing that it will poison and kill the inhabitants. But, he neither wants to kill them nor does he aim to do so, he simply wants to do his job and get paid. He views the death of the inhabitants as an unfortunate by-product of his pumping water into the cistern.

Case 12 (C12):

There is a man who gets paid for pumping water into a cistern thereby replenishing the supply of drinking water in a nearby house. Unfortunately for the inhabitants of the house, the water that the man is pumping into the cistern

to get at this root of this difference more studies would need to be run that focused on the relationship between intending to x and wanting x . It is unclear what explains why participants' responses in C6-C8 are different in this respect.

today has been systematically contaminated with a lethal poison whose effects are unnoticeable until they can no longer be cured. Even though the man pumping the water had nothing to do with poisoning the water, he knows that the water has been poisoned. Nevertheless, the man pumps the water into the cistern knowing that it will poison and kill the inhabitants. But, he doesn't care at all about the inhabitants, he simply wants to do his job and get paid.

They were then asked the five following questions:

- 1) Did the man intentionally poison the inhabitants?
- 2) Did the man want to poison the inhabitants?
- 3) Did the man intend to poison the inhabitants?
- 4) Did the man purposely poison the inhabitants?
- 5) How much blame does the man deserve for the death of the inhabitants (On a scale from 0 to 6—0 being no blame and 6 being a lot of blame)?

The results were as follows:

ANSWERS:	C11:	C12:
1) Intentionally	83%	83%
2) Wanted	5%	20%
3) Intended	50%	60%
4) Purposely	80%	88%
5) Blame	4.97	4.44

By once again following the protocol of treating analyses of intentional action as predictions about how laypersons would respond to particular cases, we have generated evidence that suggests that Anscombe's view fails as far as the folk are concerned. Hence, it appears that Miller's criticism of her water pumper scenario is supported by the pretheoretical intuitions of laypersons. Indeed, the evidence suggests that Anscombe's main contention—namely, that in order for an agent to intend to x she must desire to x—is inconsistent with the way a number of participants responded to C11 and C12. After all, even though only 13% of the participants in these two cases judged that the water pumper wanted to kill the inhabitants, 55% judged that the water pumper nevertheless

intended to do so—a highly statistically significant result ($\chi^2(1, N=80) = 32.3, p < 0.001$).

Moreover, we find the same asymmetry when we look at the participants' judgments concerning whether the water pumper intentionally killed the inhabitants. In both cases, the water pumper either explicitly does not want to kill the inhabitants or is indifferent to killing them—yet, a majority of the participants nevertheless judged that he intentionally killed the inhabitants. This result is also highly statistically significant ($\chi^2(1, N=80) = 78.6, p < 0.001$). And once again, the number of participants who judged that the water pumper brought about this side effect intentionally was greater than the number of those who judged that he intended to do so ($\chi^2(1, N=80) = 14.08, p < 0.001$). Indeed, when we add up the way all 480 participants in C1-C12 answered the first and third questions, we find a highly statistically significant difference ($\chi^2(1, N=480) = 76.8, p < 0.001$). This should finally remove any lingering doubts concerning whether the SV is consistent with ordinary usage.

Conclusion

In this paper, I have presented the results of some studies that were based on famous cases from the action theory literature and that were aimed at probing folk intuitions about desire, foresight, intention, and intentional action. Several findings were particularly noteworthy. First, people's ascriptions of intentional action are sensitive to the degree of foresight an agent has concerning the possible side effects of her actions. Other things being equal, the higher the perceived likelihood of the side effect, the more likely people are to judge that the side effect was brought about intentionally. Second,

once we compare participants' answers to the first and third questions, we find a marked asymmetry in people's ascriptions of intention and intentional action—with the latter being more likely than the former. And because the data show that non-specialists do not generally hold that intending to *x* is necessary for intentionally *x-ing*, we have reason to reject the SV as an analysis of the folk concept of intentional action. Finally, the results from C1-C12 leave little doubt that people often judge that an agent intentionally brought about a given side effect even though she did not want to do so—especially when the agent appears to be morally blameworthy for bringing it about.¹²

This last feature of the results of my studies is particularly important given that it furnishes us with additional evidence that moral considerations—especially negative ones—have a pronounced effect on how people make judgments concerning whether an agent intentionally *x-ed*. In many respects, this suggests that our actual intuitions and judgments fail to settle with philosopher's assumptions concerning the relationship between attributions of blame and ascriptions of intentional action. According to traditional rationalist models, while we may correctly rely on judgments concerning whether an action was performed intentionally in our efforts to determine someone's moral or legal responsibility, the converse is not the case—i.e., attributions of blame and praise should not affect our ascriptions of intentional action. Yet, the gathering evidence suggests that they do. Minimally, this is an interesting feature of our moral psychology that merits further research. Of course, even if future studies lend additional support to the claim that we often blame first and then gerrymander our intentionality ascriptions to

¹²Another noteworthy result is that once we compare participants' answers to the first and fourth questions we find that their intuitions concerning whether the hunter, sniper, and water pumper intentionally and purposely brought about their respective side effect line up quite nicely—51% versus 49%. This suggests that at least as far as laypersons are concerned “intentionally” and “purposely” are used synonymously—although more studies would need to be run that tested this claim specifically before we could say for sure.

justify our antecedent moral judgments—contrary to what philosophers have commonly assumed—it nevertheless remains an open question whether or not such a decision procedure is either helpful or rational. But before we can attempt to answer this question, we first need to know more about the relationship between desires, foresight, intentions, intentional actions, and moral considerations.

