

Skill, Luck, Control, and Intentional Action

Thomas Nadelhoffer

On the surface, it seems intuitively plausible that if an agent luckily manages to perform a desired action (e.g., rolling a six with a fair die or winning the lottery), the performance of which is not the result of any relevant skill on her part, we should not say that she performed the action intentionally. This intuition suggests that our concept of intentional action is sensitive to considerations of skill, luck, and causal control. Indeed, some philosophers have claimed that in order for an action to be performed intentionally it must be performed with a relevant amount of skill or control—i.e., an intentional action cannot simply be the result of luck. On this view, skill and control are necessary conditions of our everyday concept of intentional action. In this essay, I discuss empirical evidence that challenges this claim. After briefly setting the stage, I examine Al Mele and Paul Moser’s thorough analysis of intentional action—paying particular attention to some of the interesting scenarios they offer in support of their position. Next, I discuss the results of some simple psychological experiments that show that people’s judgments concerning whether actions are intentional can often be affected by the moral features of these actions—features that may trump considerations of skill, luck, and control. Finally, I conclude that if this is correct, philosophers who claim that skill and control are necessary conditions of the folk concept of intentional action appear to be mistaken.

One can test attempted philosophical analyses of intentional action partly by ascertaining whether what these analyses entail about particular actions is in line with what the majority of non-specialists would say about these actions . . . if there is a widely shared concept of intentional action, such judgments provide evidence about what the concept is, and a philosophical analysis of intentional action that is wholly unconstrained by that concept runs the risk of having nothing more than a philosophical fiction as its subject matter. (Mele, 2001, p. 27)

1. Setting the Stage

While there is wide scale disagreement among philosophers concerning how best to analyze the concept of intentional action, everyone seemingly agrees that the distinction between intentional and unintentional action plays an important role in our collective folk psychology. In a number of ordinary situations, the question of whether or not an action was performed intentionally can make a big difference in how we respond to it. As Oliver Wendell Holmes once wryly remarked, ‘even a dog distinguishes between being stumbled over and being kicked’ (Holmes, 1881/1963, p. 7). The difference between being stumbled over and being kicked is that whereas the former occurs accidentally or unknowingly, the latter is done purposely and knowingly. We react more strongly to those who harm us intentionally than to those who harm us unintentionally (or non-intentionally). If an agent neither tried to harm us nor foresaw that her actions would harm us, she is markedly less morally blameworthy for the harm caused to us than someone who either tried to harm us or recklessly brought about our harm would be.

So, for instance, if someone accidentally bumps into me on a crowded street, I readily accept her apology insofar as her bumping into me was *unintentional*. If, on the other hand, someone purposely, knowingly, and intentionally bumps into me in an effort to get even for some past action of mine, I am much less likely to excuse her behavior—especially if I believe my earlier action was justified. My different responses to these two kinds of situations suggest that at least one role that the concept of intentional action plays in our folk psychology is that of fixing blame and praise. And while we can occasionally be held morally and legally responsible for unintentional actions (e.g., cases of negligence or recklessness), intentional actions are more commonly the target of our moral judgments. But the exact nature of this role is not entirely clear.

Consequently, there is a long-standing debate in the philosophy of action concerning the nature and proper role of ascriptions of intentionality. One of the central issues of this debate is whether moral considerations affect our application of the concept of intentional action. While some scholars suggest that this concept is intimately bound up with moral considerations (e.g., Bratman, 1984; Duff, 1982, 1990; Harman, 1976; Knobe, 2003a, 2003b, 2004; Nadelhoffer, 2004, forthcoming), others claim that because the concept is used primarily to *predict* and *explain* behavior, moral considerations should not affect our ascriptions of intentional action (e.g., Butler, 1978; Katz, 1987; Mele & Sverdlik, 1996). On this latter view, 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 it not the case—i.e., attributions of blame and praise should not affect our ascriptions of intentional action.

For present purposes, I want to leave aside the *normative* question of whether moral considerations *should* act expansively on our ascriptions of intentionality and focus instead on the *empirical* question of whether they *actually* do have an effect. More specifically, I want to determine whether moral considerations sometimes

trump considerations of skill and control. If so, this would show that any analysis of the folk concept of intentional action that has skill and control as necessary conditions of intentionality would be mistaken. But we are getting ahead of ourselves. First, we should examine Mele and Moser's analysis of the folk concept of intentional action—arguably the most thorough analysis with skill and control as necessary components.

2. Mele and Moser on Intentional Action

Mele and Moser suggest that analyses of intentional action that do not make room for considerations of skill, luck, and control can be shown to 'fall prey to counter-examples' (1994, p. 223). One of their targets is Carl Ginet's claim that the following conditions are sufficient for an action's being intentional:

S's *V*-ing at *t* consists of some action, S's *A*-ing at *t*, plus that action's causing certain results or its occurring under certain circumstances, where:

- (1) S's *A*-ing at *t* was intentional, and
- (2a) at *t*, S believed of her *A*-ing that she would or might thereby *V* and,
- (2b) at *t*, in what S knew that had not slipped her mind, S had justification for the belief that was not, at the time, justification for believing a proposition too far from the truth as to how she was going to thereby *V*. (Ginet, 1990, p. 87)

In order to show that this analysis fails, Mele and Moser describe examples of actions that apparently satisfy Ginet's conditions yet are not properly said to be intentional. The first case they discuss involves an agent who luckily performs a desired action that is not the result of any relevant skill of his. It goes as follows:

Example #1 (E₁):

Consider the following case. A nuclear reactor is in danger of exploding. Fred knows that its exploding can be prevented only by shutting it down, and that it can be shut down only by punching a certain ten-digit code into a certain computer. Fred is alone in the control room. Although he knows which computer to use, he has no idea what the code is. Fred needs to think fast. He decides it would be better to type in ten digits than to do nothing. Vividly aware that the odds against typing in the correct code are astronomical, Fred decides to give it a try. He punches in the first ten digits that come into his head, in that order, believing of his doing so that he 'might thereby' shut down the reactor and prevent the explosion. What luck! He punched in the correct code, thereby preventing a nuclear explosion. (Mele & Moser, 1994, p. 224; for a similar scenario, see Mele, 1992)

By their lights, Fred's shutting down the reactor and his preventing the explosion satisfy Ginet's conditions for being intentional actions. Thus, if we adopt Ginet's analysis, we are forced to say that Fred intentionally shut down the reactor and intentionally prevented the explosion. Yet, Mele and Moser claim that this is implausible. On their view, 'when luck plays this great a role in the success of an attempt at *A*-ing, the *A*-ing is generally deemed too coincidental to count as

intentional' (1994, p. 225). To further illustrate this point, they present the following two cases:

Example #2 (E₂):

Lisa selects a sequence of six numbers to win a fair Florida instant lottery. Upon punching her six numbers into the lottery computer, Lisa wins instantly. (1994, p. 249)

Example #3 (E₃):

Mike, a normal person, is playing a game with a pair of fair dice. He will win \$20 on his next roll if and only if he throws something other than boxcars (two sixes). Mike, wanting to win, has a simple plan: he will throw a non-boxcar roll and win the money. Mike realizes that there is a slight chance that he will roll boxcars, but this does not threaten his plan. As it happens, he throws a seven. (1994, p. 252)¹

While allowing that the concepts SKILL, LUCK, and CONTROL that are under consideration are admittedly vague—much like the concept BALDNESS—Mele and Moser claim that we would ordinarily not say that these types of lucky and unskilled actions count as intentional actions (1994, p. 252). In Lisa's case, for instance, they suggest that her punching the correct numbers 'is too lucky for Lisa to have intentionally won the lottery, under the circumstances' (Mele & Moser, 1994, p. 249). Similarly, in Mike's die rolling case, Mele and Moser claim that, 'Mike lacks a kind of control over the dice required for his intentionally throwing a non-boxcar roll' (1994, p. 252). Thus, on their view, neither Lisa nor Mike performed their respective actions intentionally because in both cases the actions were not the result of any relevant skill on their part—i.e., their actions were simply the fortunate result of chance or luck.

Having provided these examples in support of their claim that skill and control are necessary conditions for intentional action, Mele and Moser proceed to offer the following analysis—an analysis they claim settles with the ordinary notion of intentional action:

Necessarily, an agent *S* intentionally performs an action *A* at time *t* if and only if:

- (1) at *t*, *S* *A*-s and her *A*-ing is an action;
- (2) at *t*, *S* suitably follows . . . an intention-embedded plan *P* of hers in *A*-ing;
- (3a) at the time of *S*'s actual involvement in *A*-ing at *t*, the process indicated with significantly preponderant probability by *S*'s on balance evidence at *t* as being at least partly constitutive of her *A*-ing at *t* does not diverge significantly from the process that is in fact constitutive of her *A*-ing at *t*; or
- (3b) *S*'s *A*-ing at *t* manifests a suitably reliable skill of *S*'s in *A*-ing in the way *S* *A*-s at *t*; and
- (4) the route to *A*-ing that *S* follows in executing her action plan *P* at *t* is—under *S*'s current circumstances—a suitably predictively reliable means of *S*'s *A*-ing at *t*, and the predictive reliability of that means depends appropriately on *S*'s having suitably reliable control over whether, given that she acts with *A*-ing as a goal, she succeeds in *A*-ing at *t*. (Mele & Moser, 1994, p. 253)

Conditions (3) and (4) are intended to prevent the three aforementioned examples of lucky or unskilled actions (i.e., (E₁)–(E₃)) from counting as intentional action, thereby fixing what Mele and Moser took to be one of the primary weaknesses of

Ginet's analysis.² The question I now want to consider is whether they are correct in claiming that their own analysis is in line with ordinary intuitions. As Mele himself points out, one way of testing attempted philosophical analyses of intentional action is to determine whether 'what these analyses entail about particular actions is in line with what the majority of non-specialists would say about these actions' (2001, p. 27). This is precisely what the following psychological studies are intended to establish.

3. Skill, Luck, Control, and Moral Considerations

Working under the rubric of 'experimental philosophy', Joshua Knobe has recently published a series of groundbreaking papers about the folk concept of intentional action (2003a,b, 2004). One of the issues that he discusses is the aforementioned question of whether the concept of intentional action operates entirely independently of evaluative considerations. In answering this question, Knobe used the four following scenarios—inspired by Mele (2001)—to conduct a few simple psychological studies:

(1A): Achievement/Skill: Jake desperately wants to win a rifle contest. He knows that he will only win the contest if he hits the bull's-eye. He raises the rifle, gets the bull's-eye in the sights, and presses the trigger. Jake is an expert marksman. His hands are steady. The gun is aimed perfectly. The bullet lands directly on the bull's-eye. Jake wins the contest.

(1B): Achievement/No Skill: Jake desperately wants to win a rifle contest. He knows that he will only win the contest if he hits the bull's-eye. He raises the rifle, gets the bull's-eye in the sights, and presses the trigger. But Jake isn't very good at using his rifle. His hand slips on the barrel of the gun, and the shot goes wild. Nonetheless, the bullet lands directly on the bull's-eye. Jake wins the contest.

(2A): Immoral/Skill: Jake desperately wants to have more money. He knows that he will inherit a lot of money when his aunt dies. One day, he sees his aunt walking by the window. Jake is an expert marksman. His hands are steady. The gun is aimed perfectly. The bullet hits her directly in the heart. She dies instantly.

(2B): Immoral/No Skill: Jake desperately wants to have more money. He knows that he will inherit a lot of money when his aunt dies. One day, he sees his aunt walking by the window. But Jake isn't very good at using his rifle. His hand slips on the barrel of the gun, and the shot goes wild. Nonetheless, the bullet hits her directly in the heart. She dies instantly. (Knobe, 2003a, p. 313)

Each subject was presented with one of these four vignettes along with the following question: 'Did Jake intentionally kill his aunt [hit the bull's-eye]?' (Knobe, 2003a, p. 313). One of the primary goals of presenting the participants with these vignettes and questions was to determine whether their judgments about these cases were affected by evaluative considerations. More specifically, Knobe wanted to ascertain whether the folk are more likely to say that Jake acted intentionally when he was attempting to perform the immoral behavior than they are when he was attempting to perform the achievement behavior.

The results of Knobe's experiments were both statistically significant and compelling. The participants' answers to the question about whether he acted intentionally were as follows:

	Immoral	Achievement
Skill	95%	79%
No Skill	76%	28%

From these results, Knobe draws the reasonable conclusion that evaluative considerations did affect the participants' ascriptions of intentional action—a result that Mele (2001) himself had explicitly predicted. Indeed, when Mele originally discussed the four scenarios that Knobe's vignettes are based upon, he acknowledged that the folk may 'treat morally significant and morally insignificant actions differently', i.e., that they might 'have a lower threshold, for example, for the intentionality of "lucky" actions deemed morally wrong than for the intentionality of equally lucky actions deemed morally neutral' (Mele, 2001, p. 38). This prediction was subsequently verified by Knobe's research. For our present purposes, Knobe's most salient finding is that moral considerations may trump considerations of skill and control when it comes to folk ascriptions of intentional action. After all, there is a very significant difference in his participants' responses to the skill cases and the no skill cases—especially when we compare the results from the no skill achievement case with those from the no skill immoral case.

However, while the data Knobe has presented are both interesting and important, there are at least three methodological problems with his vignettes. First, there is a strange tension in the no skill cases—viz., we are told both that Jake's shot *goes wild* and that his shot *hits the target*. Indeed, unless one postulates something like causal deviance (e.g., the shot goes wild but ricochets off of a tree into the target), it is hard to make sense of Jake's success. In any event, this tension may have had an effect on his participants' judgments and intuitions. Second, Jake's actions in both of the no skill cases were merely *contingently* unskilled, rather than *necessarily* unskilled; which is to say, that shooting a target—whether it be a bull's-eye or a relative—is something that an agent *could* perform skillfully even if it just so happened that in Knobe's two no skill cases, Jake's success was not a matter of skill. In order to fully test the importance of the role that skill and control play in folk ascriptions of intentional action, you would need actions that one *cannot* perform skillfully at all (e.g., rolling a six with a fair die or winning the lottery).

Finally, and most importantly, the questions that Knobe asks his participants in the achievement and immoral cases are asymmetrical. In the achievement cases, he asks the participants whether Jake *shot* the bull's-eye intentionally, whereas in the immoral cases he asks whether Jake *killed* his aunt intentionally. This difference is important because in the achievement cases, Jake was shooting the bull's-eye as a means to winning the game. In the immoral cases, on the other hand, he was shooting his aunt in order to kill her. Given that this is the case, Knobe should have asked the participants whether Jake *shot* his aunt intentionally rather than asking them whether he *killed* her intentionally. And while this may not initially appear to be a significant

problem, later in this section I will present data from some studies of mine that show that people are sometimes more inclined to judge that agents bring about the ends of their actions intentionally than they are to judge that the same agents intentionally brought about the actions that served merely as a means to their ends. In any event, these three problems with Knobe's original experiments are serious enough to weaken the conclusions he draws from his data.

So, I conducted some new experiments in an attempt to generate additional empirical support for Knobe's claim that moral considerations can sometimes trump considerations of skill, while at the same time correcting for the three aforementioned problems with his vignettes. The first of these experiments involved a close derivative of Mele and Moser's example of Fred and the nuclear reactor (E_1). Participants were 40 undergraduates, each of whom received the following vignette:

Case #1 (C_1):

A nuclear reactor is in danger of exploding. Fred knows that its exploding can only be prevented by shutting it down, and that it can be shut down only by punching a certain ten-digit code into a certain computer. Fred is alone in the control room. Although he knows which computer to use, he has no idea what the code is. Fred needs to think fast. He decides that it would be better to type in ten digits than to do nothing. Vividly aware that the odds against typing in the correct code are astronomical, Fred decides to give it a try. He punches in the first ten digits that come into his head, in that order, believing of his doing so that it 'might thereby' shut down the reactor and prevent the explosion. Amazingly, he punches in the correct code, thereby preventing a nuclear explosion and saving thousands of people.

Participants were then asked the following questions. First, did Fred intentionally punch in the correct numbers? Second, how much praise does Fred deserve for punching in the correct numbers (on a scale from 0 to 6, 0 being no praise and 6 being a lot)? Third, did Fred intentionally prevent the explosion? And finally, how much praise does Fred deserve for preventing the explosion (on a scale from 0 to 6, 0 being no praise and 6 being a lot)? The results were as follows:

- Q₁: 38% said Fred punched in the correct numbers intentionally.
- Q₂: The average praise rating was 3.0 on a 6-point scale.
- Q₃: 73% said Fred intentionally prevented the explosion.
- Q₄: The average praise rating was 4.0 on a 6-point scale.

This shows that even though most of the participants (*viz.*, 62%) were *not* willing to say that Fred intentionally punched the correct numbers, a surprising majority of them (*viz.*, 73%) *were* willing to say that Fred intentionally prevented the explosion. Thus, the participants' responses Q₃ do not line up with Mele and Moser's claim that it is implausible to think that Fred intentionally prevented the explosion. Perhaps the reason they did not expect these results is that they did not anticipate that the perceived praiseworthiness of the action would have such a marked effect on people's ascriptions of intentional action. In any event, the results of C_1 speak against Mele and Moser's analysis of intentional action—especially conditions (3) and (4)—if it is understood as an analysis of the folk concept of intentional action.

Because there is growing evidence that blame has an even more pronounced effect on ascriptions of intentional action than praise (e.g., Knobe 2003a,b, 2004; Nadelhoffer, 2004, forthcoming), I wanted to see whether a blameworthy action would similarly affect people's judgments in the Fred example. So, I conducted another survey involving a variation of C_1 . This time participants were 40 undergraduates, each of whom received the following case:

Case #2 (C_2):

Fred has just been fired from the nuclear power plant. In a desperate fit of anger, he decides to cause the reactor to meltdown. Fred knows that the only way the reactor can be forced to melt down is by punching a certain ten-digit code into a certain computer. Fred is alone in the control room. Although he knows which computer to use, he has no idea what the code is. Fred needs to think fast before the other employees return. Vividly aware that the odds against typing in the correct code are astronomical, Fred decides to give it a try. He punches in the first ten digits that come into his head, in that order, believing of his doing so that it 'might thereby' cause the reactor to meltdown. Amazingly, he punches in the correct code, thereby causing a serious nuclear meltdown and killing thousands of people.

Participants were then asked the following four questions. First, did Fred intentionally punch in the correct numbers? Second, how much blame does Fred deserve for punching in the correct numbers (on a scale from 0 to 6, 0 being no blame and 6 being a lot)? Third, did Fred intentionally cause the explosion? And finally, how much blame does Fred deserve for causing the explosion (on a scale from 0 to 6, 0 being no blame and 6 being a lot)? The results were as follows:

- Q₁: 67% said that Fred intentionally punched in the correct numbers.
- Q₂: The average blame rating was 5.23 on a 6-point scale.
- Q₃: 83% said that Fred intentionally caused the explosion.
- Q₄: The average blame rating was 5.31 on a 6-point scale.

These results show that the *blameworthiness* of Fred in C_2 had an even greater effect on participants' ascriptions of intentional action than did his *praiseworthiness* in C_1 . After all, whereas 62% of the participants in C_1 said that Fred *did not* intentionally punch in the correct numbers, 67% of the participants in C_2 said he *did* intentionally punch in the correct numbers.³ Thus, even though both praise and blame affected the subjects' judgments concerning Fred's actions, blame clearly had a more marked effect. And to the extent that in both cases Fred's punching the correct number—as well as his preventing (or causing) the explosion—was a matter of luck that did not involve any relevant skill of his, it seems clear that people's ascriptions of intentional action are more sensitive to moral considerations than they are to considerations of skill, luck, and control. This provides us with additional evidence that Mele and Moser's analysis of intentional action fails to settle with what the majority of non-specialists say.

Nevertheless, in an effort to verify that moral considerations fully explain the results of C_1 and C_2 , I ran another study. Participants were 40 undergraduates,

each of whom received the following vignette—inspired by E₂:

Case #3 (C₃):

Imagine that Fred is playing a new kind of lottery machine for \$1,000,000. In order to win, he must type in the correct ten-digit code. Vividly aware that the odds against typing in the correct code are astronomical, Fred pays his \$1 and decides to give it a try. He punches in the first ten digits that come into his head, in that order, believing of his doing so that it ‘might thereby’ win him the \$1,000,000. Amazingly, he punches in the correct code and wins the lottery!

Participants were then asked the following two questions. First, did Fred intentionally punch in the correct numbers? Second, did Fred intentionally win the lottery? The results were as follows:

Q1: 80% said that Fred did not intentionally punch in the correct numbers.

Q2: 67% said that Fred did not intentionally win the lottery.

These data obviously differ drastically from the results of C₁ and C₂. It appears that once the moral significance is stripped from the earlier cases, participants’ judgments better line up with Mele and Moser’s predictions about the implausibility of saying that unskilled or lucky actions are intentional. So even though Mele and Moser were correct to suggest that considerations of skill, luck, and control do affect folk ascriptions of intentional action, they were seemingly wrong to think that any of these considerations are necessary for an action’s being intentional. Indeed, the results of these three experiments—especially C₁ and C₂—show that when an agent’s actions are *sufficiently praiseworthy or blameworthy*, people are less interested in whether or not the action was the result of skill or luck.

Having now discussed Mele and Moser’s first two examples (i.e., E₁ and E₂), I want to briefly discuss how an earlier paper of mine sheds light on E₃—i.e., the case involving Mike and the dice game (Mele & Moser, 1994, p. 252). In Nadelhoffer (2004), I presented the results of another series of studies I ran that address the role moral evaluations play in folk ascriptions of intentional action. And while nearly all of the data from these earlier experiments are relevant to the present discussion, four surveys are particularly salient. The participants were 160 undergraduates, each of whom received one of the following vignettes and questions—inspired by the ‘Butler Problem’ (Butler 1978):

Case #4 (C₄):

Brown is playing a simple game of dice. The game requires that Brown roll a six to win. So, hoping to get a six, Brown throws a die onto the table. Unluckily for the other players, the die lands six-up and Brown wins the game.

Question: Did Brown intentionally roll a six?

Case #5 (C₅):

Same as Case #2, except the students were asked:

Question: Did Brown intentionally win the game?

Case #6 (C₆):

Brown wants to kill Smith now. Smith is in another building. There is a bomb in that building and Brown can detonate it only by producing a six-dotted image on

the lens of a camera that is focused on the top of a table in Brown's room and wired to the bomb. So, Brown takes out a normal, fair, six-sided die and tosses it onto the table, hoping that it will land six-up. Unluckily for Smith, the die lands six-up. By throwing the six, Brown detonates the bomb, thereby killing Smith.

Question: Did Brown intentionally roll a six?⁴

Case #7 (C_7):

Same as Case #5, except the participants were asked the following question:

Question: Did Brown intentionally kill Smith?

The results were as follows:

C_4 : Non-moral die case: Did Brown intentionally roll a six? Yes: 10%

C_5 : Non-moral die case: Did Brown intentionally win the game? Yes: 62.5%

C_6 : Immoral die case: Did Brown intentionally roll a six? Yes: 55%

C_7 : Immoral die case: Did Brown intentionally kill Smith? Yes: 87.5%

It seems clear that moral considerations had a startling effect on the die examples. First, whereas 87.5% of the participants of C_7 judged that Brown killed Smith intentionally, only 62.5% of the participants in C_4 judged that Brown won the game intentionally. This difference is statistically significant, $\chi^2(1, N = 40) = 6.6$, $p < 0.01$. Second, whereas 55% of the participants in C_6 judged that Brown rolled a six intentionally in the process of killing Smith, only 10% of the participants from C_4 judged that Brown intentionally rolled a six in the process of winning a game of dice. This difference is *highly* statistically significant, $\chi^2(1, N = 40) = 18.46$, $p < 0.001$.

Clearly this last result is the most puzzling of all, for how could the immorality of the outcome of Brown's rolling a six affect *whether or not he rolled the six intentionally*? In both cases, he hoped to roll the six, he had a reason for rolling six, and he rolled the die. Moreover, in both cases he lacked any control or skill over the die that made his chances of rolling a six better than one-in-six. Given that the results of C_4 and C_6 *do* reveal a statistically significant asymmetry in folk intuitions concerning whether Brown rolled a six intentionally, and provided this asymmetry *cannot* be attributed to differences in skill, it appears that the only explanation is that moral considerations do act expansively on folk judgments of intentionality. Moreover, these results lend additional support to the following two claims. First, moral considerations sometimes trump considerations of skill, luck, and control. Second, skill and control are not necessary conditions of the folk concept of intentional action as some philosophers suggest.⁵

4. Conclusion

The data I have presented in this paper give us additional reason to reject any philosophical analysis of the folk concept of intentional action that claims that in order for an action to be intentional it *must* be performed with a relevant amount of skill or control. As I suggested earlier, perhaps philosophers have mistakenly maintained that skill and control are necessary for intentional action because they underestimated the pronounced effect that moral considerations have on folk

ascriptions of intentional action. Mele, for instance, has acknowledged that people's judgments of intentional action may be sensitive to moral considerations in a way that he had not originally expected. As he asks, 'Might the folk concept of intentional action treat morally significant and morally insignificant actions differently? Might it have a lower threshold, for example, for the intentionality of "lucky" actions deemed morally wrong than for the intentionality of equally lucky actions deemed morally neutral?' (Mele, 2001, p. 38).

The growing empirical evidence suggests that this is precisely what happens. It looks like the folk concept of intentional action is more sensitive to the moral features of actions than it is to questions of skill, luck, and control. As Mele correctly points out, this complicates the 'project of capturing the folk concept of intentional action in terms of individually necessary and jointly sufficient conditions of an agent's performing an action intentionally' (2001, pp. 38–39). If nothing else, I hope to have successfully shown that this is a complication philosophers would do well not to overlook.

Before closing, however, I should make a few clarificatory remarks in order to head off some obvious—and I believe misguided—objections to the type of research that I have presented in this paper. Nowhere have I suggested that folk intuitions about intentional action will somehow solve the philosophical problems surrounding intentional action. Nor have I claimed that ordinary language or pre-theoretical judgments should serve as the *final* court of appeal. My goal has been much more modest. First, I wanted to convince the reader that folk intuitions are relevant to the debate about intentional action, and it would therefore be intellectually irresponsible to ignore them without at least giving good reason for doing so. Second, by surveying the folk, I hoped to find out which conceptual analyses of intentional action actually settle with their intuitions, so that philosophers will no longer be able to align their analyses with common sense unless their views empirically merit such support. Of course, that would not mean that analyses privileged by the endorsement of the folk are *true*, only that the burden of proof would be placed squarely on the shoulders of those who argue *contrary* to folk intuitions.

Acknowledgments

I would like to thank the following people for helpful comments on earlier drafts of this paper: Joshua Knobe, Al Mele, Eddy Nahmias, Virginia Tice, the participants at the Princeton-Rutgers Graduate Philosophy Conference, and two anonymous referees at *Philosophical Psychology*.

Notes

- [1] Mele and Moser's article is replete with interesting scenarios. However, for my present purposes I am just going to focus on these three.
- [2] Whereas condition (3) rules out unskilled actions, condition (4) rules out actions that an agent performs successfully with the requisite amount of skill, but that are nevertheless the result of luck or chance (e.g., cases of causal deviance).
- [3] $\chi^2(1, N = 40) = 6.3, p < 0.025$.

- [4] This vignette was inspired by Mele and Sverdlik's (1996) die-bombing thought experiment.
- [5] Mele has recently admitted that if people judge that Brown intentionally rolled the six in the immoral die case—and the results of C₆ show that they do—this would suggest that 'the folk concept of intentional action does not oppose intentional action to mere chance nearly as strongly as many philosophers of action do' (Mele, 2003, p. 330).

References

- Bratman, M. (1984). Two faces of intention. *Philosophical Review*, 93, 375–405.
- Butler, R. (1978). Report on *Analysis* "problem" #6. *Analysis*, 38, 113–114.
- Duff, R. A. (1982). Intention, responsibility, and double effect. *Philosophical Quarterly*, 32, 1–16.
- Duff, R. A. (1990). *Intention, agency, and criminal liability*. Oxford, England: Basil Blackwell.
- Ginet, C. (1990). *On action*. Cambridge, England: Cambridge University Press.
- Harman, G. (1976). Practical reasoning. *Review of Metaphysics*, 79, 431–463.
- Holmes, O. W. (1881/1963). *The common law*. Boston: Little, Brown and Company.
- Katz, L. (1987). *Bad acts and guilty minds*. Chicago: University of Chicago Press.
- Knobe, J. (2003a). Intentional action in folk psychology: An experimental investigation. *Philosophical Psychology*, 16, 309–324.
- Knobe, J. (2003b). Intentional action and side-effects in ordinary language. *Analysis*, 63, 190–194.
- Knobe, J. (2004). Intention, intentional action and moral considerations. *Analysis*, 64, 181–187.
- Mele, A. (1992). Recent work on intentional action. *American Philosophical Quarterly*, 29, 199–215.
- Mele, A. (2001). Acting intentionally: Probing folk notions. In B. F. Malle, L. J. Moses, & D. A. Baldwin (Eds.), *Intentions and intentionality* (pp. 27–43). Cambridge, MA: MIT Press.
- Mele, A. (2003). Intentional action: Controversies, data, and core hypotheses. *Philosophical Psychology*, 16(2), 325–340.
- Mele, A., & Moser, P. (1994). Intentional action. *Nous*, 28, 39–68. Reprinted in A. Mele (Ed.), (1997). *The philosophy of action*. New York: Oxford University Press.
- Mele, A., & Sverdlik, S. (1996). Intention, intentional action, and moral responsibility. *Philosophical Studies*, 82, 265–287.
- Nadelhoffer, T. (2004). The Butler problem revisited. *Analysis*, 64, 277–283.
- Nadelhoffer, T. (forthcoming). Praise, side effects, and folk ascriptions of intentional action. *Journal of Theoretical and Philosophical Psychology*.