

# Intentional Action, Folk Judgments, and Stories: Sorting Things Out

ALFRED R. MELE AND FIERY CUSHMAN

How are our actions sorted into those that are intentional and those that are not? The philosophical and psychological literature on this topic is livelier now than ever, and we seek to make a contribution to it here. Our guiding question in this article is easy to state and hard to answer: How do various factors—specifically, features of vignettes—that contribute to majority folk judgments that an action is or is not intentional interact in producing the judgment? In pursuing this question we draw on a number of empirical studies, including some of our own, and we sketch some future studies that would shed light on our topic. We emphasize that the factors that concern us here are limited to features of stories to which subject respond: examples include the value of the action asked about, the agent's being indifferent to performing that action, and the agent's seeking to perform it. We do not discuss underlying cognitive or emotional processes here, nor do we discuss whether respondents are making errors of any kind. (Both of these issues are discussed in Cushman and Mele [forthcoming].)

## 1. THREE KINDS OF ACTION

In the present section we draw some distinctions that set the stage for our discussion of empirical results. Our actions have effects, and an agent's *bringing about* such an effect is itself an action. For example, unbeknownst to Ann, her unlocking the door to her house frightened an intruder. That is, at least one effect of Ann's unlocking her door was the intruder's fright. Her bringing about this effect—that is, her frightening the intruder—is an action. *Side-effect actions*, as we understand this

expression, are defined partly in terms of the effects of actions that the agent seeks to perform and succeeds in performing. The following will do as a definition:  $X$  is a side-effect action performed by an agent  $S$  if and only if  $S$  successfully seeks to perform an action  $A$ ,  $E$  is an effect of his so doing,  $X$  is his bringing about  $E$ , and  $X$  has the following properties:  $S$  is not at the relevant time seeking to  $X$  either as an end or as a means to an end, and  $X$  is not in fact a means to an end that  $S$  is seeking at the relevant time. (Some things we do are means to ends we are seeking at the time even if we do not conceive of them as such then. For example, Cal, who has no concept of calories, is running in order to lose weight. Because Cal's bringing it about that he burns more calories than usual is a means to his losing weight, we do not count it as a side-effect action—even though Cal has no idea that his bringing this about is a means to his end. Hence, the final clause in our definition of *side-effect action*.) Given this definition, Ann's frightening the intruder is a side-effect action. Ann successfully seeks to unlock her door, the intruder's fright is an effect of her so doing, her frightening the intruder is her bringing about that effect, and frightening the intruder is neither something that she was seeking to do at the time (either as an end or as a means to an end) nor a means to an end that she was seeking in acting.

Although we are interested in side-effect actions in general, side-effect actions that are *anticipated* by their agents—*ASEAs*, for short—are of special interest to us. There are at least two ways for an action not to be an *ASEA*: (i) the agent may seek to perform it (either as an end or as a means); (ii) the agent may neither seek to perform it nor anticipate performing it. We dub actions of the first sort *seek-type actions* and actions of the second sort *0-type actions*. Actions of all three types are discussed below.

## 2. PROBABILITY IN SEEK-TYPE ACTIONS

Some seek-type actions are judged not to be intentional by folk majorities. The following story by Joshua Knobe provides a case in point:

JAKE (N). Jake desperately wants to win the 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 his sights, and presses the trigger. But Jake is not 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. (Knobe 2003b, 313; see Mele 1992, 363–71 for a similar story and for discussion.)

Readers who deem it false that Jake intentionally hit the target—as seventy-two percent of Knobe's respondents did—may attempt to explain why by appealing, for example, to Jake's lack of a relevant skill or the associated low probability that his attempt would be successful.<sup>1</sup> The influence of factors of this kind on folk

1. Some may appeal instead to the causal deviance in the story. On this, see section 3.

judgments about whether or not a seek-type action is intentional is the topic of the present section.

Our study 1 included a story very similar to JAKE (N):

RIFLE. Lydia desperately wants to win the rifle contest. She has taken one class on shooting, and the instructor told her again and again how she had a natural talent that was as good as any expert. Therefore, Lydia is just as confident that she will hit the bull's-eye as experts are about their own chances of hitting bull's-eyes. Of course, without any experience, she is not very good at using a rifle—her instructor was just trying to be nice. She raises the rifle, gets the bull's-eye in the sights, and pulls the trigger. Her hand slips on the barrel of the gun, and the shot goes wild. Nonetheless, the bullet hits the bull's-eye. Lydia wins the contest.

Did Lydia intentionally hit the bull's-eye?

In study 1, our 150 respondents answered sixteen questions of the form “Did *S* intentionally *A*?” on a seven-point scale: “1” was a strong “no” and “7” a strong “yes.”<sup>2</sup> The mean response to RIFLE was 4.23. Counting answers of 5 or higher as “yes” answers and answers of 3 or lower as “no” answers, fifty-two percent of our respondents answered yes and forty-five percent answered no.<sup>3</sup>

We also tested a related story:

POOL. Tom is on a first date at a pool hall. He wants to impress his date by hitting a difficult shot, the red striped ball in the side pocket. He takes a close look and aligns the cue stick for the necessary shot. But Tom knows that he is not a very good pool player. His hand is shaking and he does not guide the cue stick properly. Nevertheless, the cue stick hits the ball just right and the shot is a success: the red striped ball goes in the side pocket. Tom's date is very impressed.

Did Tom intentionally hit the red striped ball into the side pocket?

The mean response here is much higher: 5.41 (with seventy-five percent saying yes and twenty-one percent saying no). Why might that be? Perhaps Tom is seen as having more control over the success of his attempt than Lydia has over the success of hers.<sup>4</sup> Possibly, Tom's task seems easier than Lydia's and, unlike hers, is seen as something that a novice has a reasonable chance of succeeding at. Also, although Tom has difficulty handling the cue stick and Lydia has difficulty handling the rifle, there is an extra element of luck in Lydia's story: she hits the bull's-eye even though “the shot goes wild.” A lucky ricochet must have been involved. This contributes to the impression that Lydia has less control over the success of her attempt than Tom

2. All sixteen stories appear in the appendix of Cushman and Mele (forthcoming).

3. Given how similar RIFLE is to JAKE (N), the difference between our results and Knobe's is puzzling. Does Lydia's *confidence* make the difference?

4. On the place of such control in intentional action, see Mele and Moser (1994).

has over the success of his. (On the assumption that Lydia's plan was to hit the bull's-eye directly, some causal deviance was involved in her hitting it. The impact of causal deviance on "intentionally" judgments is discussed in section 3.)

Notice that although Tom's subjective probability that he will succeed in his attempt in POOL apparently is low, as it should be, Lydia's subjective probability that she will succeed in her attempt in RIFLE is high. Even so, POOL gets a much higher "intentionally" rating than RIFLE. This suggests the hypothesis that when people are asked whether seek-type actions are intentional, they treat the agent's subjective probability of success as less important than the objective probability of success.

Another story in study 1 (based on Mele 2001, 33–34) bears on this hypothesis:

**BOWL.** Earl is an excellent and powerful bowler. His friends tell him that the bowling pins on lane 12 are special 200-pound metal pins disguised to look like normal pins for the purposes of a certain practical joke. They also tell him that it is very unlikely that a bowled ball can knock over such pins. Apparently as an afterthought, they challenge Earl to knock over the pins on lane 12 with a bowled ball and offer him ten dollars for doing so. Earl believes that his chance of knocking over the pins on lane 12 is very slim, but he wants to knock them down very much. He rolls an old bowling ball as hard as he can at the pins, hoping that he will knock down at least one. To his great surprise, he knocks them all down! The joke, it turns out, was on Earl: The pins on lane 12 were normal wooden ones.

Did Earl intentionally knock down the pins?

Earl's subjective probability that he will knock down some pins is very low, but the objective probability of his knocking some down is very high. The mean "intentionally" rating for BOWL was 6.36 (the highest for the sixteen stories). This provides some confirmation for the current hypothesis. Notice also that this hypothesis is linked to our suggestion that POOL gets a significantly higher "intentionally" rating than RIFLE partly because Tom's task seems considerably easier than Lydia's. Other things being equal, the easier the task, the more likely the agent is to succeed at it.

Regarding neutral seek-type actions, the objective probability of success seems to have a substantial effect on folk ascriptions of intentionality. The effect is not matched in cases of good and bad seek-type actions, as we will explain.

Knobe's study of shaky Jake included a bad seek-type action:

**JAKE (B).** 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. He raises his rifle, gets her in his sights, and presses the trigger. But Jake is not 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 2003b, 313)<sup>5</sup>

Whereas only twenty-eight percent of Knobe's respondents to the neutral story about Jake said that Jake intentionally hit the bull's-eye, seventy-six percent of his respondents to the story about the bad action said that Jake intentionally killed his aunt (Knobe 2003b, 313–14).

Knobe also tested a story in which an agent luckily succeeds in doing something good (Knobe 2003b, 320). In RADIO, Klaus tries to save many innocent lives by disabling a communications device. He "raises his rifle, gets the device in his sights, and presses the trigger. But Klaus is not very good at using his rifle. His hand slips . . . and the shot goes wild. Nonetheless, the bullet lands directly in the communications device. The mission is foiled and many innocent lives are saved." Knobe reports that ninety-two percent of his thirty-eight subjects said that Klaus intentionally hit the device.

One neutral story in study 1 (DICE) is about a person who wanted to win a dice game, knew that he would do so if and only if he rolled an eleven on his next toss of a pair of dice, and rolled an eleven. The mean response to the question whether he intentionally rolled an eleven was 2.41 (with nineteen percent of respondents saying yes and seventy-three percent saying no). Thomas Nadelhoffer (2004, 281) tested the following "bad" analogue of this story (based on Mele and Sverdlik 1996, 279; also see Mele 2003, 329–30 for discussion):

**KILLING DIE.** 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. By throwing a six, Brown detonates the bomb, thereby killing Smith.

A total of 87.5 percent of his forty respondents answered "yes" to the question "Did Brown intentionally kill Smith?", while fifty-five percent of another group of forty respondents answered "yes" to the question "Did Brown intentionally roll a six?".

The results for the following neutral version of the story were different:

**GAME DIE.** 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. (Nadelhoffer 2004, 279)

5. For discussion of a pair of cases similar to JAKE (N) and JAKE (B), see Mele (2001, 39–40).

Sixty-two percent of Nadelhoffer's forty respondents answered "yes" to the question "Did Brown intentionally win the game?", and ten percent of another group of forty respondents answered "yes" to the question "Did Brown intentionally roll a six?".

The difference in the responses about whether Brown intentionally rolled a six in Nadelhoffer's pair of stories is interesting. One might have thought that almost everyone would see *intentionally* rolling a six with a single toss of a fair die as something that requires having special powers over dice—powers ordinary people do not have. In light of this, the ten percent figure for those who say that Brown intentionally rolls a six in GAME DIE (like the nineteen percent figure for a similar toss in DICE) might have seemed high, and the jump to fifty-five percent in the response to KILLING DIE is striking. The salient difference between the two stories, of course, is a difference in *value*—the difference between the neutral goal and outcome in the former story and the bad goal and outcome in the latter.

One may be tempted to draw the following conclusion about seek-type actions from the data reported in this section: Although most people have objective-probability standards for intentional action in the case of neutral seek-type actions, they have no objective-probability standards at all for intentional action in the case of valenced seek-type actions. However, this would be hasty. Notice that in all of the valenced stories discussed in this section, lives hang in the balance. It may be that although objective-probability standards for intentional action are extremely low or even nonexistent in life-and-death cases of this kind, they are significantly higher for valenced seek-type actions in which significantly less is at stake. It would be interesting to see, for example, whether or not, as badness gradually diminishes in luck-featuring stories in which someone successfully attempts to do something bad, so does the "intentionally" rating. What would happen if a headache were substituted for a death in KILLING DIE? (Brown wants to give Smith a headache, and by throwing a six he releases gas that has the desired result. The headache, which is moderately painful, lasts for an hour.) A story of this kind is worth testing.

A claim that *is* supported by the data reviewed in this section is that folk majorities make "intentionally" judgments as though, other things being equal, they have significantly higher objective-probability standards for neutral intentional seek-type actions than for good or bad intentional seek-type actions.<sup>6</sup> Of course, when "other things being equal" clauses are not empty, at least one other factor is relevant. We take up such a factor in the following section.

### 3. CAUSAL DEVIANCE

Philosophers have discussed deviantly caused actions in connection with the project of providing an analysis of intentional action. Study 1 included two stories highlighting causal deviance, BEE (based on Mele 1987) and WEEDS:

6. Notice our use of "as though" here. It may be that the goodness and badness of seek-type actions have an effect that overrides assessments made on the basis of objective probability. As we mentioned, we avoid discussion of underlying cognitive and emotional processes in this article.

BEE. Mark wants to provide the right answer to a multiple-choice question. He thinks that the right answer is “bee,” which is option “c.” But he is rushing, and he circles the letter “b” instead of the letter “c” next to the word “bee.” As it happens, “bee” was not the right answer, but “b” (“ant”) was. Luckily for Mark, the answer he circled was correct.

Did Mark intentionally provide the right answer?

WEEDS. Jen sees some bothersome weeds growing next to her driveway and wants to eliminate them. She decides to go to the hardware store to buy weed spray. As she pulls out of her garage, the wheel slips in her hand and she drives off to the side of the driveway. All the weeds are crushed and killed under the car. With the weeds eliminated, Jen doesn’t need to go to the hardware store.

Did Jen intentionally eliminate the weeds?

The mean “intentionally” ratings for the two stories were very low, the lowest for the sixteen stories: BEE 1.36, WEEDS 1.67.

Our RIFLE (reproduced above) also involves causal deviance, though the deviance is not as salient as it is in BEE and WEEDS. In RIFLE, Lydia hits the bull’s-eye even though “the shot goes wild.” So the causal path from the pulling of the trigger to the hitting of the bull’s-eye included something important that diverged from what one would assume to be Lydia’s plan. Presumably, she planned to hit the bull’s-eye directly; but, in fact, she must have benefited from a ricochet. Again, the mean response to the question whether Lydia intentionally hit the bull’s-eye was 4.23. Is the difference between this middling mean response and the strong negative mean responses to BEE and WEEDS explained by a difference in the salience of the deviance or by something else?

Bertram Malle (2006) tested a version of Knobe’s JAKE (N) (on which RIFLE is based) in which the deviance was made salient. The clause about the wild shot reads as follows: “the shot goes wild, bouncing off a heavy post” (99). Seventy percent of his respondents judged that Jake did not hit the bull’s-eye intentionally. As we mentioned, counting answers of 5 or higher as “yes” answers and answers of 3 or lower as “no” answers, fifty-two percent of our respondents to RIFLE answered yes and forty-five percent answered no. Perhaps the greater salience of the deviance in Malle’s story accounts for the difference between his results and ours. But the “not intentional” response is even stronger in BEE and WEEDS than it is in Malle’s version of JAKE (N). Over ninety-five percent of the respondents to BEE answered no, and less than three percent answered yes; and the figures for WEEDS were similar: ninety percent no, less than seven percent yes. Even when the deviance in a RIFLE-type story is made salient, it does not have as strong an effect as the deviance in BEE and WEEDS does.

This fact is interesting. As we mentioned, Knobe finds an asymmetry in responses to JAKE (N) and JAKE (B). Whereas only twenty-eight percent say that Jake intentionally hit the bull’s-eye, seventy-six percent say that he intentionally killed his aunt. And Malle got similar results while making the deviance salient: eighty-five percent for Jake’s shooting his aunt intentionally versus thirty percent

for hitting the bull's-eye intentionally (Malle 2006, 99). When cases of deviantly produced neutral actions have much lower “intentionally” ratings than these bull's-eye cases, as BEE and WEEDS do, would the “intentionally” ratings for bad—or good—counterpart actions also be significantly lower than the ratings Knobe and Malle got for their versions of JAKE (B) and the rating Knobe got (in RADIO) for the man, Klaus, who saved many innocent lives with a lucky, causally deviant shot? This is worth testing.

Our study 2 was motivated partly by this question. It included the following variants of BEE and WEEDS in which a *bad* action is performed:

**BOMB.** Mark wants to detonate a bomb to kill some innocent people. He knows that only one button will detonate the bomb and that the other buttons will deactivate it. He thinks that the right button has the initials “R.E.D.” on it. But he is rushing, and he presses the red colored button instead. As it happens, the red colored button was the button for detonating the bomb. Mark detonates the bomb and kills some innocent people. Did Mark intentionally kill some people?

**PUPS.** Jen sees some puppies playing next to her driveway again and wants to kill them. She decides to go to the hardware store to buy some rat poison that she thinks will work on the puppies. As she pulls out of her garage, the wheel slips in her hand and she drives off to the side of the driveway. All the puppies are crushed and killed under the car. With the puppies eliminated, Jen doesn't need to go to the hardware store. Did Jen intentionally kill the puppies?

Respondents were volunteers in basic undergraduate philosophy courses at Florida State University. Whereas the mean response for BOMB was 5.12 ( $N = 57$ ), that for PUPS was much lower: 2.68 ( $N = 58$ ). Why might that be? Is this difference in responses explained by the difference in degree of harm caused, by the difference in the kinds of deviance, or by a combination of the latter differences?

This is testable. A variant of BOMB in which puppies are substituted for innocent people may be tested. And here is a pertinent variant of PUPS:

**BOYS.** Jen sees some small boys playing next to her driveway again and wants to kill them. She decides to go to the gun store and buy a rifle to kill them with. As she pulls out of her garage, the wheel slips in her hand and she drives off to the side of the driveway. The boys are crushed and killed under her car. With the boys eliminated, Jen doesn't need to go to the gun store. Did Jen intentionally kill the boys?

Testing stories like these would provide more evidence about whether, other things being equal, people treat causal deviance as a much greater obstacle to a seek-type action's being intentional when the action is neutral than when it is good or bad.

#### 4. MISTAKES

Sometimes people forget to do what they intended to do. For example, they intended to buy some milk on the way home from work, but they forget and drive home, as usual. In situations of this kind, what do people do intentionally and what do they do unintentionally? Scenarios of this kind call for attention.

Some actions are not easily classified as either side-effect actions or seek-type actions. Consider the following vignette:

DRIVE. Almost every day, Al drives home from work along the same two-mile route. He drives east on Store Street for a mile and turns left onto Home Street. His house is a mile north on Home. While eating lunch in his office today, Al decides to drive to his favorite store later in the day to pick up some beer and chips before driving home. He wants to have them on hand for a game he plans to watch on TV. The store is east on Store Street a half mile beyond Home Street. Al's plan is to drive east on Store to the store, buy the beer and chips, and then drive home. But he forgets. When he gets to Home Street, he turns left, as usual. Shortly after he gets home, he remembers what he had planned to do. A little upset with himself, he gets back in his car and drives to the store.

In study 3, fifty undergraduates at Florida State University read DRIVE. Half were asked whether Al intentionally turned left, and half were asked whether he intentionally drove the rest of the way home. Seventy-two percent of each group answered no. These results suggest that "mistaken" actions of the kind at issue are generally regarded as not intentional.

It would be interesting to see whether *valenced* analogues of DRIVE elicit significantly different "intentionally" judgments. Here are two relevant sketches of stories. In DRIVE (B), Al promised to pick up a friend at the airport after work, and the stranded friend is robbed (or worse) by someone who offered to drive him home. Did Al intentionally drive home? In DRIVE (G), Al forgets to do the same thing and then, upon arriving home, rescues his next-door neighbor from an attempted robbery (or worse). (The friend takes a taxi home without mishap.) Did Al intentionally drive home?

#### 5. SIDE-EFFECT ACTIONS, VALUE, AND INDIFFERENCE

Like seek-type actions, side-effect actions may be good, bad, or neutral. Any value-neutral analysis of intentional action that counts only seek-type actions as intentional (Malle and Knobe 1997; Mele and Moser 1994) will count all side-effect actions as nonintentional, no matter which of these three groups they fall into. For the most part, the results of pertinent studies are consistent with such an analysis in the spheres of good and neutral side-effect actions anticipated by the agent (good and neutral *ASEAs*, for short). But in the sphere of bad *ASEAs*, things apparently change dramatically. In this section, we explore various features of stories that may account for the asymmetry.

Here is a well-known vignette by Knobe:

CHAIRMAN (HARM). The vice-president of a company went to the chairman of the board and said, “We are thinking of starting a new program. It will help us increase profits, but it will also harm the environment.” The chairman . . . answered, “I don’t care at all about harming the environment. I just want to make as much profit as I can. Let’s start the new program.” They started the new program. Sure enough, the environment was harmed. (Knobe 2003a, 191)

Knobe reports that eighty-two percent of his thirty-eight respondents said that the chairman intentionally harmed the environment (Knobe 2003a, 192). In our study 1, the mean answer for the question whether the chairman intentionally harmed the environment was 5.79, with eighty-four percent of our respondents saying yes.

What accounts for the high “intentionally” rating? Candidates for answers include, among other things, the *badness* of harming the environment, the chairman’s *indifference* to harming it, his *confidence* that starting the program would harm it, and various interactions among the preceding candidates.

We start with the *badness* of the *ASEA*—something that is contrasted, obviously, both with neutrality and with goodness. Study 1 included two stories about neutral *ASEAs* (BIRDS and FATTY below). The mean “intentionally” rating for the neutral actions was 3.085 (“not intentionally”). Regarding good *ASEAs*, Knobe found that when helping the environment replaced harming it in his chairman story, there was a dramatic reversal of opinion: seventy-seven percent of the forty respondents said that the indifferent chairman did *not* intentionally help the environment (Knobe 2003a, 192). And we found, similarly, that the mean “intentionally” rating for this good *ASEA* in our 150 respondents was 2.36, with seventy-nine percent of our respondents saying no. The badness of the chairman’s harming the environment certainly seems to be playing a role.<sup>7</sup>

Does the chairman’s *indifference* to harming the environment help to account for the high “intentionally” rating? Our study included the following vignette about an agent who is *not indifferent* to a bad *ASEA*:

POND. Al said to Ann: “You know, if you fill in that pond in the empty lot next to your house, you’re going to make the kids who look for frogs there sad.” Ann replied: “I know that I’ll make those kids sad. I like those kids, and I’ll definitely regret making them sad. But the pond is a breeding ground for mosquitoes; and because I own the lot, I am responsible for it. It must be filled in.” Ann filled in the pond, and, sure enough, the kids were sad.

Did Ann intentionally make the kids sad?

7. We are not taking a stand on whether the badness plays a direct role or an indirect role that is mediated by its contribution to the blameworthiness of the agent (Phelan and Sarkissian, forthcoming).

The mean “intentionally” rating was 3.19 (with only twenty-eight percent saying yes)—much lower than our results for CHAIRMAN (HARM). However, because the lack of indifference to the bad *ASEA* is not the only difference between POND and CHAIRMAN (HARM), one should not conclude that this difference fully accounts for the difference in responses. For example, harming the environment may strike people as considerably worse than making kids sad; and Ann may strike people as justified in filling in the pond at the expense of the kids’ sadness, whereas the chairman may strike people as unjustified in starting the profit-making venture at the expense of harming the environment. We return to this observation shortly.

Mark Phelan and Hagop Sarkissian (forthcoming) tested a story (PLANNER) that is similar to POND in that the agent of a bad *ASEA* displays proper concern rather than indifference. In PLANNER, a city planner replies as follows when he is informed that implementing a plan for cleaning up toxic waste would “also increase the levels of joblessness”: “I feel terrible about increasing joblessness. But we have to do something about our pollution problem. Let’s start the project.” Implementing the plan had both of the predicted results. Like Ann in POND, the planner is not indifferent about the bad *ASEA*. Phelan and Sarkissian report that only twenty-nine percent of their twenty-one respondents judged that the planner “increased levels of joblessness intentionally.” These results are similar to ours for POND.

Obviously, lack of indifference to the negative *ASEA* is not the only feature that distinguishes PLANNER from CHAIRMAN (HARM). For example, the decision to clean up toxic waste at the expense of increasing unemployment may strike people as justified whereas the chairman’s decision strikes them as unjustified. Elsewhere, one of us suggested testing a version of the “harm” story in which the chairman responds along the following lines to the information he receives: “I truly wish that I could make money for this company without harming the environment. Unfortunately, that seems to be impossible. Reluctantly, I’m instructing you to start the new program” (Mele 2006, 287). Here we suggest testing a version of POND in which Ann reports that she does not care at all about making the kids sad. (The same can be done with a variant of PLANNER, of course.) Data on these stories would provide evidence about the effect of an agent’s indifference to a bad *ASEA*.

Does the chairman’s *confidence* that the program will harm the environment help to account for the relatively high “intentionally” rating? A study by Nadelhoffer provides relevant evidence. He tested a trio of stories about a hunter who wants to shoot a deer and is indifferent to the prospect of shooting a bystander. In HUNTER 1, “The hunter realizes that if he shoots the deer, the bullet will *definitely* hit” a bystander (Nadelhoffer 2006, 142). (Presumably, the bullet will pass through the deer and strike the bystander.) In HUNTER 2, “could *possibly*” replaces “will *definitely*.” And HUNTER 3 features “a very small chance” (143). In each case, the hunter “does not care at all” about the bystanders. Each of Nadelhoffer’s three groups of forty respondents saw one story only. Whereas sixty-eight percent of the respondents to the first story said that the hunter intentionally shot the bystander, this was said only by thirty percent of the respondents to the second story and thirty-five percent of the respondents to the third.

Nadelhoffer also tested a trio of stories about a sniper who has been ordered to kill an enemy commander and is indifferent to alerting “the other enemy soldiers to his presence” with his gunshot (2006, 147). Now, a sniper’s alerting the enemy to his presence might put him in grave danger, and that might be a very bad thing from the sniper’s point of view. However, because the respondents are informed that “the sniper doesn’t care at all about” alerting the enemy, they might not see his alerting the enemy as dangerous or in any way bad for him: they may see this action as a *neutral* side-effect action. In any case, when the featured term was “definitely,” fifty-five percent said that the sniper intentionally alerted “the enemies to his presence,” as compared with forty percent for “probably” and twenty percent for “a very small chance.” (There were forty respondents to each story.)

Nadelhoffer’s data on his hunter stories suggest that if the chairman’s accurate subjective probability of harming the environment were considerably lower, his harming the environment would get a significantly lower “intentionally” rating. This is worth testing.

Notice that we said “accurate subjective probability.” In Nadelhoffer’s stories, the agent “realizes” that thus and such will definitely (or possibly, probably, etc.) happen if he *A*-s: that is, the agent *believes* this *and* what he believes is *true*. These studies leave all of the following possibilities open: in making “intentionally” judgments about bad side-effect actions, people are equally sensitive to the agent’s subjective probabilities and the actual probabilities (or the probabilities a reasonable person would assign); they are significantly more sensitive to the former; they are significantly more sensitive to the latter. To get evidence about this, we tested the following story in study 2.

HUNTER (U). A hunter is about to shoot at a deer with a high-powered rifle. It doesn’t occur to him that if he shoots the deer, the bullet will pass through the animal and hit the bird watcher he sees standing just beyond it, even though he certainly should have known that this would happen. He shoots and the bullet hits the bird watcher.

Did the hunter intentionally shoot the bird watcher?

The hunter’s shooting the bird watcher is what, in section 1, we called a *θ*-type action: the hunter neither seeks to shoot the bird watcher nor anticipates doing so. The mean response to our question was 2.79 (N = 59), with nineteen percent answering yes. The primary difference between this story and Nadelhoffer’s HUNTER 1 is that although his hunter realizes that he will definitely hit the bystander, ours should realize this but does not. The objective probability is the same. Even so, whereas sixty-eight percent of Nadelhoffer’s respondents answered yes to the “intentionally” question, only nineteen percent of ours did. Apparently, high *subjective* probability plays an important role in motivating people to see bad side-effect actions as intentional. (Notice the contrast here with seek-type actions; see section 2.) But what would happen if the agent’s high subjective probability were much higher than the objective probability (and he performs the *ASEA*)? Is it *accurate* high subjective probability that is doing the work, or high subjective probability alone? This is another question that calls for tests. One promising test

case features a version of Nadelhoffer's HUNTER 1 in which the hunter *falsely* believes that "if he shoots the deer, the bullet will *definitely* hit" a bird watcher and, in fact, there is only a slight chance of this. In the imagined story, despite the remote chance of the bullet passing through the deer, the hunter hits the deer and the bird watcher with the same bullet.

Earlier, we asked about the effect of the chairman's indifference in the "harm" scenario. What about his indifference to helping the environment in the "help" scenario? What effect does it have? One way of eliminating the indifference would amount to placing the chairman's helping the environment outside the sphere of side-effect actions. If, in instituting the program, he were motivated both by a desire to make a profit and by a desire to help the environment, his helping the environment would be a seek-type action—not a side-effect action. But there is another way of eliminating his indifference. Elsewhere, one of us suggested testing a variant of Knobe's "help" story in which the chairman is represented as honestly saying something along the following lines: "I am extremely pleased that I will make money for this company while also helping the environment. But, in all honesty, I must tell you that the terms of my employment do not allow me to take the environment into account in making policy decisions. I would instruct you to start this new program even if I did not believe that it would help the environment" (Mele 2006, 287). In this variant, the chairman is not seeking to help the environment even though he is not indifferent to its being helped. Data generated by a story like this would provide evidence about the effect the chairman's indifference to helping the environment has on respondents to Knobe's story.

Here is a story from our study 1 in which an agent who performs a good *ASEA* is not indifferent to it:

COACH. Coach Don decided during mid-season that his son Donnie should become a starting player on their high-school baseball team: Donnie had become one of the best players on the team. Don knows that Donnie's being a starting player would make both Donnie and himself happy. Like anyone, Don likes to be happy, and he likes his son to be happy. But Don takes his responsibility as a coach very seriously. He does not allow himself to consider personal happiness when making decisions about his ball team. Don starts Donnie in the next game, and they both feel happy.  
Did Don intentionally make Donnie happy?

In our story, Don treats his son's happiness—a good thing—as a side effect of his action of making him a starting player, but he is not indifferent to his son's happiness. Among our 150 respondents, the mean "intentionally" rating for COACH was 3.93, whereas that for CHAIRMAN (HELP) was 2.36. Whereas the rating for COACH places it neither in the "intentional" nor in the "not intentional" category, the rating for CHAIRMAN (HELP) is firmly in the latter category.

The salient difference between Knobe's "harm" and "help" stories about the chairman is the difference between a bad effect and a good one—the environ-

ment's being harmed versus its being helped. And, as we have reported, the stories yield asymmetrical majority judgments about whether the featured *ASEA* is intentional. Does an agent's indifference to a good *ASEA* have a stronger effect than indifference to a neutral *ASEA* would have?

One source of data about neutral *ASEAs* is a study by Nadelhoffer (2006). He used three vignettes. One (*EAGLE*) is similar to his hunter stories, the difference being that the hunter realizes that if he shoots the deer he will definitely cause a beautiful eagle to fly away. The others are similar to his sniper stories. In *AIR*, the sniper realizes that the bullet he fires at the enemy commander will definitely move air molecules, and he is indifferent about that. And in *HOT*, heating the barrel of his gun replaces moving air molecules. Each story had forty respondents. The results are as follows: thirty-five percent said that the hunter intentionally caused the eagle to fly away, forty-five percent said that the sniper intentionally disturbed some air molecules, and sixty-eight percent said that the sniper intentionally heated the barrel of his gun.

In study 1, as we mentioned, we tested two stories about neutral *ASEAs*:

**BIRDS.** Stacey is very fond of cardinals—she thinks they are a particularly lovely bird. She goes to the garden center and gets a bird feeder that is designed to attract lots of cardinals. An employee at the garden center tells Stacey that the feeder is also sure to attract blue jays. “I don’t care about the blue jays,” says Stacey. “As long as it attracts cardinals, I am happy.” Stacey sets up the feeder, and it attracts cardinals and blue jays. Did Stacey intentionally attract blue jays to her back yard?

**FATTY.** Susan’s friend works as a chef in a diner. Her friend is sick and asks Susan to take over for a few days. Susan is disgusted by all the grease and fatty foods in the diner. In fact, she is so disgusted that she decides to stop eating fatty foods. When she tells this to her nutritionist, he replies that the change in diet will be perfectly healthy, but will cause Susan to lose some weight. Susan replies that losing weight doesn’t matter one way or the other to her. She just thinks fat and grease are disgusting. Susan stops eating fatty foods, and she loses some weight. Did Susan intentionally lose some weight?

The mean response was 3.31 for *BIRDS* and 2.86 for *FATTY*. Thirty-three percent of our respondents answered yes about Stacey and twenty-six percent answered yes about Susan.

We asked whether an agent's indifference to a good *ASEA* has a stronger effect than indifference to a neutral *ASEA*. Here is a summary, in table form, of the percentage results we reported for the six stories we discussed in this connection—*CHAIRMAN (HELP)* and a quintet of stories about *ASEAs*:

Good and Neutral *ASEAs* with Agent Indifference

Story	“intentionally”
<b>Chairman (help)</b>	<b>17% (ours); 23% (Knobe’s)</b>
Air	45%
Birds	33%
Eagle	35%
Fatty	26%
Hot	68%

The results suggest that indifference to a good *ASEA* does more to motivate “not intentionally” judgments than indifference to a neutral *ASEA* does.<sup>8</sup> (Incidentally, given the premium placed on thinness in American culture, many respondents may see Susan’s losing weight in *FATTY* as good rather than neutral.)

Earlier we noted that an agent’s high *subjective* probability that he will *A* (as compared with high objective probability of which the agent should have been cognizant but was not) plays an important role in motivating people to see *bad* side-effect actions as intentional. What about *neutral* side-effect actions? In study 2, we tested the following version of *HOT*:

*HOT* (U). A hunter is about to shoot at a deer with a high-powered rifle. It doesn’t occur to him that shooting will cause the barrel of his gun to get hot, even though he certainly should have known that this would happen. He shoots and the barrel heats up.

Did the hunter intentionally heat the barrel of his gun?

The mean response to this question about a *0*-type action was 3.48 ( $N = 58$ ), with thirty-four percent answering yes—half of the percentage figure for a yes answer to Nadelhoffer’s *HOT*. Here again, the difference is striking. And once again, we would like to know whether it is the combination of high subjective probability and high objective probability—that is, *accurate* high subjective probability—that produces the effect or high subjective probability alone.

In this section, we explored various features of stories that may account for an observed asymmetry in folk “intentionally” judgments about bad *ASEAs*, on the one hand, and good and neutral *ASEAs*, on the other. The badness, goodness, and neutrality seem to have distinguishable effects. Indifference and confidence also seem to have effects. We suggested some further tests that would shed light on interactions among these factors.

8. Nadelhoffer’s *HOT* is the outlier in the group of five neutral *ASEA* stories here. It is difficult to know what to make of that. If more studies of this kind were done, perhaps a stable, noteworthy pattern would emerge. In conversation about the results for Nadelhoffer’s trio of neutral *ASEA* stories, Randy Clarke suggested that respondents may be sensitive to differences in “causal closeness.” The thought is that, other things being equal, the nearer in a causal chain a neutral anticipated side effect is seen by subjects as being to the intentional action that causes it, the stronger the intuition will be that the side-effect action of generating it is an intentional action. This is testable.

## 6. SOME NOTABLE ASYMMETRIES

Knobe's chairman stories elicit asymmetrical "intentionally" responses. Although most people say that the chairman who helps the environment does not intentionally help it, most people say that the chairman who harms the environment intentionally harms it.<sup>9</sup> Here the asymmetry apparently is associated with what might be termed a good-versus-bad asymmetry in the stories. However, there also is a good-versus-bad asymmetry in the RADIO and JAKE (B) pair of stories; and even so, most people count the actions at issue in these stories as intentional. What accounts for this difference? The difference between *ASEAs* and seek-type actions plainly makes a difference here. The good and bad actions featured in the chairman stories—harming and helping the environment—are *ASEAs*, whereas the good and bad actions featured in RADIO (disabling the enemy's communication device) and JAKE (B) (killing his aunt) are seek-type actions.

Positive and negative valences of actions seem to pull in the same direction when the intentionality of seek-type actions is being assessed and in opposite directions when the intentionality of *ASEAs* to which the agents are indifferent is the issue. This effect may be associated with intuitive moral assessments of the *agents* themselves for their conduct in the stories. The fact that an agent is seeking something bad tends to say something negative about his character, and so does the fact that an agent is indifferent to a bad *ASEA* of his. However, although the fact that an agent is seeking something good tends to say something positive about his character, the fact that an agent is indifferent to a good *ASEA* of his does not: indeed, it says something negative about him. It may be that positive and negative valences assigned to *agents* for their conduct (e.g., "bad" for the chairman in the help scenario) have a greater effect on folk "intentionally" judgments than do the positive and negative valences of the *actions* that subjects are asked to assess as intentional or not intentional (e.g., "good" for helping the environment).

Another difference between reactions to *ASEAs* and reactions to seek-type actions also merits emphasis. In the case of the *ASEAs* we reviewed (with the exception of Nadelhoffer's HOT), good *and neutral* actions are lumped together as not intentional. So there seems to be, more fully, a good-or-neutral-versus-bad asymmetry in folk judgments about whether *ASEAs* are intentional. However, regarding some seek-type actions with a low objective probability of success, most people count the good *and bad* actions as intentional and the parallel neutral actions as not intentional. (Compare the results for RADIO and JAKE (B) with those for JAKE (N), and compare the results we reported for tossing a six or an

9. As we explain in Cushman and Mele (forthcoming), some respondents may see the chairman's harming the environment as a seek-type action rather than an *ASEA*. In a study reported there, we made it clear that harming the environment was a side-effect action by having the foreseen harm start to occur ten years after profits were made for the current year's balance sheet. The mean response for those who saw just this story was 4.51—well below our mean response of 5.79 to CHAIRMAN (HARM) in study 1. Other respondents also read a parallel "help" story. Those who responded to the help story first displayed a mean response of 3.59 to our harm story (and 2.20 to our help story), and those who responded in the reverse order displayed a mean response of 4.14 to the harm story (and 2.48 to the help story).

eleven intentionally with one another.) Here there is a neutral-versus-good-or-bad asymmetry. Also, in the case of *ASEAs*, the combination of foresight that by doing what one plans to do one will do something bad and indifference about that is treated as though it is sufficient for the *ASEA*'s being intentional, whereas this is not so when what the agent is indifferent about is a good or neutral *ASEA*. The difference between *ASEAs* and seek-type actions is an important one in folk "intentionally" judgments.

## 7. CONCLUSION

Researchers and theorists have lots of different, legitimate aims in their investigations of folk judgments about what is done intentionally and what is not. Some are in search of the folk concept of intentional action, others are looking for two or more distinct folk concepts of intentional action, yet others are concerned to identify the underlying psychological processes that produce such judgments, and so on. Our aim in this article was to shed light on various features of stories that have a significant effect on folk "intentionally" judgments and on interactions among these features and to sketch new studies that would shed additional light on these matters. Because the data at the heart of the experimental investigation of the folk concept(s) of intentional action are lay responses to stories, our project here bears directly on each of the aims identified.<sup>10</sup>

## REFERENCES

- Cushman, F., and Mele, A. Forthcoming. "Intentional Action: Two-and-A-Half Folk Concepts?" In *Experimental Philosophy*, ed. J. Knobe and S. Nichols. New York: Oxford University Press.
- Knobe, J. 2003a. "Intentional Action and Side-Effects in Ordinary Language." *Analysis* 63:190–94.
- . 2003b. "Intentional Action in Folk Psychology: An Experimental Investigation." *Philosophical Psychology* 16: 309–24.
- Malle, B. 2006. "Intentionality, Morality, and Their Relationship in Human Judgment." *Journal of Cognition and Culture* 6: 87–112.
- Malle, B., and Knobe, J. 1997. "The Folk Concept of Intentionality." *Journal of Experimental Social Psychology* 33: 101–21.
- Mele, A. 1987. "Intentional Action and Wayward Causal Chains: The Problem of Tertiary Waywardness." *Philosophical Studies* 51: 55–60.
- . 1992. "Acting for Reasons and Acting Intentionally." *Pacific Philosophical Quarterly* 73: 355–74.
- . 2001. "Acting Intentionally: Probing Folk Notions." In *Intentions and Intentionality: Foundations of Social Cognition*, ed. B. Malle, L. Moses, and D. Baldwin, 27–43. Cambridge, MA: MIT Press.
- . 2003. "Intentional Action: Controversies, Data, and Core Hypotheses." *Philosophical Psychology* 16: 325–40.
- . 2006. "The Folk Concept of Intentional Action: A Commentary." *Journal of Cognition and Culture* 6: 277–90.
- Mele, A., and Moser, P. 1994. "Intentional Action." *Noûs* 28: 39–68.

10. We are grateful to Adam Feltz and Karen Foulke for assistance with the administration of the FSU surveys and data management.

- Mele, A., and Sverdlik, S. 1996. "Intention, Intentional Action, and Moral Responsibility." *Philosophical Studies* 82: 265–87.
- Nadelhoffer, T. 2004. "The Butler Problem Revisited." *Analysis* 64: 277–84.
- . 2006. "Desire, Foresight, Intentions, and Intentional Actions: Probing Folk Intuitions." *Journal of Cognition and Culture* 6: 133–57.
- Phelan, M., and Sarkissian, H. Forthcoming. "The Folk Strike Back; or Why You Didn't Do It Intentionally, Though It Was Bad and You Knew It." *Philosophical Studies*.