

**Enough Skill to Kill:  
Intentionality Judgments and the Moral Valence of Action**

Steve Guglielmo<sup>1</sup>     Bertram F. Malle  
University of Oregon

Extant models of moral judgment assume that an action's intentionality directs assignments of blame. Knobe (2003b) challenged this fundamental order and proposed instead that the blameworthiness of an action directs (and thus unduly biases) people's intentionality judgments. His and other researchers' studies suggested that morally significant actions are seen as intentional even when the agent lacks skill (e.g., killing somebody with a lucky shot) whereas equivalent neutral actions are not (e.g., luckily hitting the bull's eye). The present six studies offer an alternative account of these provocative findings. We show that in previous studies the morally significant action (e.g., killing) and the corresponding neutral action were not equivalent. In particular, killing — but not hitting the bull's-eye — can be accomplished by a basic action (pressing the trigger) for which even an unskilled agent has sufficient skill. When this basic action is performed unskillfully or is absent, however, people less often view the killing as intentional (Studies 1 and 2). Moreover, hitting the bull's-eye is more difficult than killing, and difficult actions are less often judged intentional (Studies 3-5). When difficulty across actions is held constant, intentionality judgments are highly sensitive to skill information regardless of moral valence (Study 5 and 6). The present studies thus disconfirm the hypothesis of a moral evaluation bias in intentionality judgments and instead document people's sensitivity to subtle features of human action.

Some observations about moral judgment are uncontroversial. Foreseeability has been widely observed as a requirement for folk responsibility judgments — people normally do not consider someone responsible for outcomes that the person was unable to anticipate (Shaver, 1985). Equally significant, Hamilton (1978) recognized the role of obligation, as people are blamed only for negative outcomes that they were obligated to prevent. Finally, Weiner (1995) observed that such obligation is meaningful only if the outcome was controllable by the person, if he or she could have intentionally prevented it. Thus, blame for negative events arises when the person *could have* and *should have* prevented it (Malle, Moses & Baldwin, 2001).

These findings highlight people's rational, rule-following assignments of responsibility and blame (cf. Nichols & Mallon, 2006). Admittedly, things are not always so clear-cut. Alicke (2000) and others have shown that extraneous variables can alter moral judgment beyond rational rules. But there is no doubt that such rules exist and operate in many instances.

One of the most important rules regarding the assignment of blame is that intentional moral transgressions — when the person intentionally performs the immoral action — amplify blame (Heider, 1958; Shaver, 1985; Ohtsubo, 2007). Intentional harmful actions are by definition foreseeable and controllable and blatantly violate the obligation not to cause harm. Actually having a desire and intention to bring about harm and exerting effort to realize this

---

<sup>1</sup> Please address correspondence to Steve Guglielmo via email: [sgugliel@uoregon.edu](mailto:sgugliel@uoregon.edu)

intention is the worst offense in any social community. Excuses sometimes avert blame for unintentional harm; only justifications can possibly avert blame for intentional harm (Tedeschi & Reiss, 1981), and except for rare cases in which a compelling reason for the harmful act is available (e.g., the dentist hurting her patient), full blame applies.

Judgments of intentionality, thus, serve as central inputs to judgments of blame. Schematically, the social perceiver takes two judgment steps, in order:

1. Determine: Is the negative event intentional?
2. If Yes → examine the actual intention or goal; then assign proper blame.  
If No → examine controllability (preventability), obligation, and foreseeability; then assign proper blame.

Knobe (2003a, b) challenged this fundamental order and proposed instead that the blameworthiness of an action can bias people's intentionality judgments. In particular, Knobe argued, the same behavior that is seen as unintentional when performed without moral implication (e.g., shooting and hitting a bull's eye) may in fact be seen as intentional when performed with moral implication (e.g., shooting and killing another person). If true, such a pattern would cast serious doubt on rational models of blame assignment. People would not, as traditionally believed, assess intentionality to designate blame but would instead assess blame to designate intentionality.

The fault of the traditional account of blame, following Knobe's argument, lies in its assumption about how intentionality judgments are made. Both philosophical theories (e.g., Mele, 1992; Mele & Moser, 1994; Mele & Sverdlik, 1996) and psychological theories of intentionality (Malle & Knobe, 1997) subscribe to the valence-neutral model sketched in Figure 1. According to this model, for positive, negative, and neutral behaviors alike, people process five information components that all have to be present for a behavior to be considered intentional. If even one component is missing, the behavior does not count as intentional.




Figure 1. A model of the folk concept of intentionality.

Knobe claims that this five-component model is correct only for neutral actions but does not hold for negatively valenced actions (we will return to the question of positive actions later in this paper). Knobe's evidence for his claim falls into two sets. The first set challenges the necessity of the *intention* component for judgments of intentionality. According to the standard intentionality model, a behavior is judged as intentional only if the agent actually intended to perform that behavior (Adams, 1986; Malle & Knobe, 1997). In contrast, Knobe (2003a) provided data suggesting that people consider a known but *unintended* side-effect intentional if that effect is negatively valenced (e.g., harming the environment, risking the lives of soldiers, decreasing sales). Other researchers have replicated this effect (Cushman & Mele, 2008; McCann, 2005; Nadelhoffer, 2006a; Nichols & Ulatowski, 2007), and we refer to this set of findings as the "side-effect challenge."

The second set of findings questions the necessity of the *skill* component for judgments of intentionality. According to the standard intentionality model, a behavior is judged as intentional only if the agent has reliable ability or skill to produce that behavior (Malle & Knobe, 1997; Mele & Moser, 1994; Thompson, Armstrong, & Thomas, 1998). Knobe (2003b) showed that an agent's unskilled neutral action (e.g., a lucky shot to win a contest) is not viewed as intentional but an equivalent unskilled *immoral* action (e.g., a lucky shot to kill someone) is very much seen as intentional. Once more, other researchers have replicated this effect (Nadelhoffer, 2004, 2005; Sousa & Holbrook, under review), and we refer to this set of findings as the "skill challenge."

In a separate paper, we have analyzed the validity of the side-effect challenge (Guglielmo & Malle, under review), particularly its conditions of occurrence and the components of intentionality it reveals. We came to the conclusion that, once comparisons between morally valenced and nonvalenced cases are made truly parallel, and once people can express their judgments in their own terms, the side-effect finding disappears. People are in fact keenly sensitive to the basic components of intentionality that the standard model postulates — belief, desire, and intention — and do not seem to be biased by the moral valence of a side-effect. Moreover, people do not consider known but unintended side-effects intentional, and so they do not make judgments of intentionality without a prior judgment of intention.

The skill challenge, however, still looms. To examine this challenge we will try to identify the conditions under which such findings occur and clarify the findings' implications for theories of intentionality, theories of blame, and for the prospect of valid *mens rea* judgments in the law. For if the negative valence of a defendant's action biases jurors toward "seeing" intent in the action, we would seriously question their capacity to assess such intent.

### *Knobe's Original Study*

Knobe's (2003b) original skill challenge derived from a study in which four components of intentionality (see Fig. 1) were held constant but skill was varied. Together with a manipulation of valence, Knobe employed a 2 (skill: high/low)  $\times$  2 (valence: neutral/negative) design. The specific vignettes were as follows:

#### Neutral:

Jake desperately wants to win the rifle contest. He knows that he will only win the contest if he hits the bulls-eye. He raises the rifle, gets the bull's-eye in the sights, and presses the trigger.

[skill:] Jake is an expert marksman. His hands are steady. The gun is aimed perfectly...

[no skill:] 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.

Did Jake intentionally hit the bull's-eye?

#### Negative:

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 the sights, and presses the trigger.

[skill:] Jake is an expert marksman. His hands are steady. The gun is aimed perfectly...

[no skill:] 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.

Did Jake intentionally kill his aunt?

As expected, when skill was high, both valence conditions elicited intentionality judgments in most people (79% for neutral, 95% for negative). When skill was absent, however, the neutral condition showed the predicted drop in intentionality judgments to 28% (because, according to the standard model, a necessary condition of intentionality was missing) whereas the negative condition showed persistently high intentionality judgments by most people (76%). This last result — that people judged the action to be intentional even when the (seemingly necessary) skill component of intentionality was missing — constitutes the skill challenge.

The force of the skill challenge rests on the assumption that the neutral and the negative conditions differ only in valence and that the negative valence keeps intentionality judgments high despite the absence of a core component of intentionality, namely skill. We will raise two doubts about this assumed equivalence between conditions and thus about the hypothesis that valence can bias intentionality judgments. First, we will try to show that even in the low-skill negative condition, the perpetrator performed a basic action for which he had sufficient skill, and that this action may count as intentional killing (Studies 1 and 2). Second, we will try to show that the two actions in question (killing vs. hitting the bull’s eye) are not equally difficult and that a given amount of skill is therefore sufficient to perform one of them intentionally but not the other (Studies 3 to 5). Once differences in skill and difficulty are held constant, we then suggest, any asymmetries in intentionality judgments due to valence disappear (Study 6).

Our strategy is to account specifically for the negative/no skill condition, the core of the skill challenge. A neutral condition can usually be omitted because people respond normatively to it, following the standard model of intentionality (see Malle & Knobe, 1997). Also, the negative/high-skill condition has the obvious result of leading to almost uniform intentionality judgments. The puzzling results occur specifically in the negative/no skill condition, where we see frequent intentionality judgments despite the apparent absence of skill. If we can identify the critical factors that keep intentionality judgments high in this case, we may be able to meet Knobe’s challenge.

## Study 1

In Knobe’s (2003b) original study, people judged a wobbly, wayward shot as intentional when it fulfilled the plan to kill another person. Why would people interpret this action as intentional even if it lacked skill?

Recall the description of the agent’s low skill in Knobe’s vignette: “[Jake] raises his rifle, gets her in the 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...” Although the agent is described as “not very good,” his lack of skill does not manifest until *after* he has pressed the trigger. So the act of pressing the trigger is indisputably intentional and itself requires little skill. According to Mele (2003), intentionally performing any action (such as pulling a trigger) that leads to a morally charged outcome (e.g., someone’s death) may be sufficient for people to deem the outcome intentional. Slipping on the barrel (and thus displaying lack of skill) may thus be irrelevant to people if the crucial intentional action of pressing the trigger has already been skillfully performed. Further, if pressing the trigger counts as killing, then killing is skillful (and intentional) as well. It follows that if the initial act of pressing the trigger were portrayed as lacking skill, killing should lack skill as well and intentionality judgments should drop. Study 1 therefore described a variant of the shooting vignette in which the agent slips *before* the act of pulling the trigger, which would raise doubts about the intentionality of pulling the trigger and, by hypothesis, lower the rate of intentionality judgments for killing the aunt.

### Method

Participants were 175 undergraduate students who completed a larger computer-presented questionnaire and received partial course credit in return. Each was randomly assigned to one of two conditions, which differed in whether the agent slips before or after pulling the trigger. In the “slip after” condition (identical to Knobe’s original vignette), the agent “raises his rifle, gets her in the sights, and pulls 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.” In contrast, in the “slip before” condition, Jake’s “hand slips on the barrel of the gun even before he pulls the trigger.” To provide a tougher test, this new “slip before” condition omitted the sentence that describes Jake’s dispositional lack of skill (“Jake isn’t very good at using his rifle”). Thus, only the slipping before pulling the trigger cast doubt on the agent’s skill, and only slipping (not, for example, a complex interaction between slipping and having low skill) could reduce intentionality judgments. In both conditions, the bullet hits the aunt directly in the heart and she dies instantly. All participants then answered with Yes or No the standard intentionality question, “Did Jake kill his aunt intentionally?”<sup>2</sup>

### Results and Discussion

When Jake slipped after he pulled the trigger, 93% said that the act of killing was intentional; when he slipped before he pulled the trigger, 71% said the killing was intentional, even though they were not told outright that Jake lacked skill in the latter condition. A loglinear analysis showed this difference to be significant,  $z = 3.8$ ,  $p < .001$ ,  $d = .61$ . As predicted, even when the outcome is held constant (and is highly negative), people are notably less likely to call the action intentional when the agent’s lack of skill becomes evident earlier in the causal chain of events and has a causal effect on the action. These findings suggest that one reason for the high intentionality ratings in Knobe’s negative/no skill condition (and the “slip after” condition in this study) was that the “unskilled” agent successfully performed a critical action that led to the killing — namely, pressing the trigger. The agent had enough skill to perform this action intentionally, and to the extent that pulling the trigger counts as killing, the agent had enough skill to intentionally kill as well. However, when the slip occurred *before* pulling the trigger, the agent’s skill was in doubt and fewer people felt that the killing was intentional.

But why did 71% of people still ascribe intentionality to the killing in the “slip before” condition? Importantly, this condition did not state that Jake *failed* to press the trigger, only that the slippage occurred before pressing the trigger. Participants may have believed that the agent corrected himself after slipping and still successfully pressed the trigger. Therefore, if the agent failed to press the trigger altogether, people should be even less likely to view the killing as intentional. Study 2 tested this prediction.

## Study 2

### Method

The current vignette — unlike Knobe’s and those in Study 1 — did not explicitly mention the agent pulling the trigger. The key sentences in this new “no pull” version were: “He raised his rifle and got her in the sights, but his hand slipped on the barrel of the gun and a shot

---

<sup>2</sup> The exact formulation of all vignette materials used in the current studies can be viewed online at <http://www.uoregon.edu/~sgugliel/materials/skill>

went off. Bouncing off a nearby post, the bullet nonetheless hit his aunt directly in the heart. She died instantly.” Just as the “slip before” condition from Study 1, the “no pull” condition omitted reference to Jake’s dispositional lack of skill (“Jake isn’t very good at using his rifle”) to ensure that only the absence of pulling the trigger could reduce intentionality judgments. We predicted that people in this “no pull” condition would be less likely to view the killing as intentional than those in the original Knobe study and either scenario from Study 1 (“slip before” or “slip after”).

Participants were 50 undergraduate students who completed a larger computer-presented questionnaire and received partial course credit in return. Each read the no-pull vignette, then answered the intentionality question: “Did James intentionally try to kill his aunt?” and “Did James kill his aunt intentionally?” In addition, people answered an intention (trying) question, “Did James intentionally try to kill his aunt?” This latter question was included to ensure that our manipulation truly weakened the skill condition, not the perceived strength of the intention to kill.

### *Results*

As predicted, intentionality judgments for killing were substantially lower in the no-pull condition (42%) than in Study 1’s slip-after condition (93%) and slip-before condition (71%). In a loglinear analysis, these differences were highly significant,  $z_s = 6.3$  and  $2.9$  (both  $ps < .01$ ),  $d_s = 1.06$  and  $0.56$ , respectively. Despite this tendency not to view the act of killing as intentional, most people (86%) still felt that the agent intentionally *tried* to kill his aunt. Thus, people clearly recognized the agent’s prior negative intention, but they did not automatically judge the performed action to be intentional.

### *Discussion*

These results support the hypothesis that people’s intentionality judgments are sensitive to variations in skill. In particular, when the constitutive action of pulling the trigger was no longer skillfully performed, the overarching action of killing was infrequently seen as intentional. Thus, one reason for the high intentionality ratings for the “unskilled” killing in Knobe’s initial study was that the killing was not actually perceived to be unskilled. Because the protagonist in Knobe’s study performed the act of pressing the trigger with sufficient skill (and thus did so intentionally) and because pressing the trigger counts as killing, that act of killing was seen (by most) as intentional as well. When, in our studies, the act of pressing the trigger became unskilled (Study 1) or was absent (Study 2), the killing was no longer constituted by a skillful action and its perceived intentionality dropped substantially.

We still have a puzzling finding to address, however. In Knobe’s neutral condition, where only 28% of people saw the act of hitting the bull’s-eye intentional, the agent also pressed the trigger intentionally: “He raises the rifle, gets the bull’s-eye in the sights, and presses the trigger... The bullet lands directly on the bull’s-eye. Jake wins the contest.” If intentionally pressing the trigger counts as intentionally killing the aunt in the negative condition, then why does it not count as intentionally hitting the bull’s-eye in the neutral condition?

In order to answer this question, we need to consider the interplay between skill and difficulty (Heider, 1958; Weiner, 1986). In the first two studies, we showed that, given a certain level of action difficulty, people’s intentionality judgments for that action vary with the agent’s skill. If skill is low (e.g., because the person slips before pulling the trigger), people are substantially less likely to see the action as intentional. The following studies reverse this logic and examine whether, given a certain level of the agent’s skill, people’s judgments of

intentionality vary with the action's difficulty. If action difficulty is high, we would expect people to be less likely to see the action as intentional. Studies 3 and 4 vary the difficulty of the agent's negative action and assess its effect on judgments of intentionality. Study 5 then tests directly whether difficulty, not moral valence, accounts for the difference between Knobe's (2003b) immoral and neutral conditions.

### Study 3

Knobe's (2003b) claim of a moral bias in intentionality judgments relied on the comparison between intentionally killing a person and intentionally hitting the bull's-eye. Besides moral significance, however, there is a difference between the two conditions in the implied size of the target the agent hits — and thus the difficulty of the action. Whereas “hitting the bull's-eye” specifies a precise condition for success (having the bullet land directly on the small bull's-eye), “killing the aunt” offers a greater number of success conditions — for example, Jake could hit many different body parts to succeed in killing his aunt. Some of these body parts will be more similar in size to a bull's eye whereas others will not. Study 3 examined whether the size of the body part that a shooter hits influences people's intentionality judgments. Given whatever level of skill the protagonist has, the smaller (and the more similar to the bull's-eye) the body part is, the more difficult it is to hit it, which should decrease people's inclination to view the hit as intentional.

#### *Method*

Participants were 98 undergraduate students who completed a larger computer-presented questionnaire and received partial course credit in return. To test the hypothesis that size of body part influences intentionality judgments, we created three different scenarios in which the protagonist shot the target in either the arm, the hip, or the head. In addition, we varied the protagonist's skill at shooting, using the same skill/no-skill manipulation as in Knobe's (2003b) study. Therefore, the study had a 3 (body part)  $\times$  2 (skill level) between-subjects design.

We chose the three body parts to approximate a scale of target size from larger (arm) to medium (hip) to smaller (head). We predicted that intentionality judgments should decrease from arm to head, corresponding to a decrease in target size (and an increase in action difficulty). But we also chose these particular body parts to simultaneously test Knobe's competing hypothesis that blameworthiness drives intentionality judgments. People should view hitting the head as the most injurious and blameworthy (and therefore, according to Knobe, the most intentional), followed by hitting the hip, and finally hitting an arm, which should be the least blameworthy and therefore the least intentional. To absolutely ensure that the condition in which the protagonist hit the target's head was the most blameworthy, we added to the story that the target died instantly in this case but omitted this sentence in the other two cases.

The two accounts thus make competing predictions:

Action difficulty account: Intentionality judgments for Arm > Hip > Head.

Valence account: Intentionality judgments for Head > Hip > Arm.

The vignette that tested these predictions described a protagonist Jason who lived in a remote Alaskan community and had a dispute with his neighbor over a stretch of land. The critical portions of the story, with the manipulations indicated accordingly, were as follows:

One day in 1999 Jason decided to teach his neighbor a lesson. He waited until he spotted the neighbor walking onto the disputed land, raised his rifle, and set the sights on the neighbor's right arm [hip] [head].

[no skill:] But Jason was a terrible shot, so as he pulled the trigger, his hand slipped, and the shot went wild. Jason did not see the flight of the bullet, but as a matter of fact, it bounced off a large boulder several feet away and subsequently hit the neighbor's right arm [hip] [head. The neighbor died instantly].

[skill:] Jason was an expert marksman, so he pulled the trigger and directly hit the neighbor's right arm [hip] [head. The neighbor died instantly].

Each participant received one of six possible vignettes, then answered both an intentionality question (“Did Jason intentionally shoot his neighbor in the right arm [hip] [head]?”) and a blame question (“How much blame does Jason deserve for shooting his neighbor?” on a 0-6 scale).

### *Results*

As predicted by both competing hypotheses, people provided more blame when the agent hit the neighbor in the head and killed him ( $M = 5.3$ ) than when he hit the man either in the arm or hip ( $M = 4.7$ ),  $F(1,92) = 6.9$ ,  $p < .05$ . More important, as predicted by the target size account but contradicting the valence account, a loglinear analysis revealed a linear trend in intentionality judgments across body parts such that hitting the arm was most often judged as intentional (94%), followed by hitting the hip (80%), followed by hitting the head (69%),  $z = 2.1$ ,  $p < .05$ . This differentiation tended to be stronger for the low-skill agent (for whom task difficulty presumably has more of an impact), but the body part  $\times$  skill interaction term did not reach significance. Similarly, the overall main effect for skill was noticeable — people viewed the skilled agent's behavior as more intentional (88%) than the non-skilled agent's behavior (74%) — but did not reach significance.

### *Discussion*

Although people provided the highest blame judgments when the agent shot his neighbor in the head, people were least likely to view this action as intentional, suggesting that blame judgments were not driving intentionality judgments. In fact, these two judgments were uncorrelated ( $r = .07$ , *ns*), contrary to what Knobe's model of valence-biased intentionality would suggest. Instead, target size (a proxy for difficulty) predicted variations in intentionality judgments: hitting a smaller target (the victim's head) was less likely to be seen as intentional than hitting a larger target (the victim's arm).

## Study 4

To replicate the pattern of findings in Study 3 and to gather a direct measure of perceptions of difficulty, we conducted a second study with a similar vignette content as in Study 3, but focusing on the critical conditions of low-skill agents.

### *Method*

Participants were 100 undergraduate students who completed a larger computer-presented questionnaire and received partial course credit in return. This time the agent “decided to teach his neighbor a lesson and shoot him in the body [collarbone].” Participants answered an intentionality question (“Did Frank intentionally hit his neighbor's body [collarbone]?”), a difficulty question (“How difficult was it for Frank to shoot the neighbor's body [collarbone]?” on a 0-7 scale), and a blame question (“How much blame does Frank deserve for hitting his neighbor's body [collarbone]?” on a 0-7 scale).

## Results

As predicted, 80% of people said that hitting the body was intentional, whereas 64% said that hitting the collarbone was intentional,  $z(98) = 1.74$ ,  $p$  (one-tailed)  $< .05$ ,  $d = .36$ . Moreover, people judged the body shot to be far less difficult ( $M = 3.3$ ) than the collarbone shot ( $M = 4.9$ ),  $F(1, 98) = 13.9$ ,  $p < .001$ ,  $d = .74$ . Blame judgments, however, did not follow a parallel pattern, being very similar for the body shot ( $M = 6.2$ ) and the collarbone shot ( $M = 6.4$ ). Thus, the target size manipulation was connected to the difficulty of the action ( $r = -0.35$ ,  $p < .001$ ) and marginally to its intentionality ( $r = 0.18$ ,  $p = .08$ ), but not to its blameworthiness ( $r = -0.06$ ,  $p > .50$ ).

## Discussion

Studies 3 and 4 showed that intentionality judgments vary with the difficulty of the action — here operationalized as the size of the shooting target. Given a certain level of skill, relatively easy actions (e.g., shooting a whole arm or body) are more often seen as intentional than relatively difficult actions (e.g., shooting at the head or collarbone), irrespective of their blameworthiness. We thus hypothesize that the act of killing in Knobe's (2003b) original story was seen as relatively easy, even for the unskilled protagonist, because it only required hitting the victim *somewhere* lethal. By contrast, the act of hitting the bull's-eye was seen as relatively difficult because there was only one way of accomplishing it.

Study 5 directly tested the hypothesis that killing is normally judged as easier than hitting the bulls-eye. At the same time, Study 5 also attempted to equate the two actions in difficulty. This dual aim was met by asking people to make two intentionality judgments: the standard one about the act of killing and a new one about the act of shooting the aunt in the heart. If people's intentionality judgments are sensitive to the interplay of skill and difficulty, a substantial number of people should see the relatively easy act of killing as intentional but the more difficult shooting in the heart as unintentional. The consequence of the aunt's death is of course constant, so any differences between the two intentionality judgments are independent of blameworthiness and can only be attributed to features of the action itself. Our predicted results would be consistent with other findings on the skill challenge. For example, Nadelhoffer (2004, 2005) reported that when an agent rolled a six on a die to detonate a bomb and kill his enemy, people more often viewed the more general act of killing as intentional than they did the specific act of rolling a six.

## Study 5

### Method

Participants were 120 undergraduate students who completed a larger computer-presented questionnaire and received partial course credit in return. Two participants had missing values and were excluded. Each read either the low-skill negative or low-skill neutral vignettes from Knobe's (2003b) original study, then answered a difficulty question ("How challenging was the shot that James took?") and a skill question ("How much skill do you think James had when he took the shot?"), both on a 0-6 scale. Participants in the negative condition then answered a specific intentionality question ("Did James hit his aunt in the heart intentionally?") and the standard general intentionality question ("Did James kill his aunt intentionally?"). To keep the conditions as similar as possible, participants in the neutral condition also answered a specific intentionality question ("Did James hit the bull's-eye intentionally?") followed by a more general one ("Did James win the contest intentionally?").

## Results

Skill judgments, as expected, were quite low overall and did not differ notably between the neutral ( $M = 1.57$ ) and negative conditions ( $M = 1.22$ ),  $t(118) = 1.59$ ,  $p > .10$ ,  $d = .30$ . Difficulty, however, varied by valence conditions. People viewed shooting the bull's-eye as more difficult ( $M = 3.78$ ) than shooting the aunt ( $M = 2.83$ ),  $t(116) = 2.79$ ,  $p < .01$ ,  $d = .51$ .

The pattern of intentionality judgments showed a striking pattern. Within the negative condition, the general question about killing elicited almost perfect intentionality rates (98%) whereas the specific question about hitting the aunt in the heart elicited very low intentionality rates (27%),  $z(58) = 6.4$ ,  $p < .001$ ,  $d = 1.64$ .<sup>3</sup> Thus, the identical action elicited remarkably different intentionality judgments depending on whether it was described as “killing” or “hitting in the heart.”

Within the neutral condition, the question about winning elicited moderate intentionality judgments (43%), and the question about hitting the bull's-eye did not differ from this rate (38%),  $p > .10$ ,  $d = .10$ . This suggests that winning and hitting the bull's-eye were of comparable specificity. After all, there was only one way of winning the contest — by hitting the bull's eye.

Finally, we compared intentionality rates for the two specific (and comparably difficult) acts — the neutral one of hitting the bull's-eye and the negative one of hitting the aunt's heart — and they did not differ,  $z(118) = -1.30$ ,  $p > .19$ ,  $d = .24$ . If anything, the negative one was slightly lower (27% vs. 38%).

## Discussion

The results of this study suggest that Knobe (2003b) and other researchers confounded difficulty with valence when they used the verb phrases *killing the aunt* and *hitting the bull's eye*. With this choice of verbs, people saw the negative action of killing as easier to accomplish than the neutral one of hitting the bull's eye. For an agent with low skill, an easier action has a greater chance of being performed intentionally because even a low level of skill may be sufficient for successful performance. Thus, the standard negative condition in past studies yielded more intentionality judgments than the neutral condition because the former was easier, not because it was morally significant. When we correct for this confound and equate the specificity and difficulty of the actions at issue — comparing hitting the aunt's heart with hitting the bull's eye — intentionality judgments for the negative action drop substantially. Only 27% of participants deemed the negative action of hitting the aunt's heart intentional, as low a rate as the neutral condition in Knobe's (2003b) original demonstration and its various replications (Nadelhoffer, 2004, 2005).

The studies so far provide considerable evidence for the claim that Knobe's finding of a “skill neglect” in intentionality judgments arose not because of special moral considerations but because of two factors: (a) the availability of a basic action that was performed intentionally (i.e., pressing the trigger) and (b) differential levels of difficulty of the actions in question (killing vs. hitting the bull's eye). By staying within the standard vignette, however, we have not been able to completely equate the respective difficulty of the morally significant action and the control

---

<sup>3</sup> These within-subject dichotomous data can be analyzed in a variety of ways, from nonparametric tests to repeated-measures ANOVA (assuming that underlying the dichotomous intentionality judgment is a continuous state of confidence). All analyses lead to the same results, and we are reporting here the more conservative Wilcoxon's signed rank test and its corresponding effect-size estimate.

action. We might try several different verbs and may get closer to such equality, but the two actions differ in too many other features. Our last study thus examined an action that was the same in the control and the experimental conditions.

In addition, we wanted to clarify an ambiguity surrounding the valence of the control action of hitting the bull's eye. Knobe (2003b) described it as an achievement and asked people to rate how much praise it deserves. However, he also pointed out that an achievement is not *morally* praiseworthy (which is why we have characterized it as morally neutral in our studies). In his fourth study, Knobe (2003b) examined intentionality judgments for a truly morally praiseworthy act — aiming at a military device, hitting it, and thereby saving a number of lives. In this study, 92% of people found even the “unskilled” shot that saved lives intentional, just as many as found the unskilled shot that killed the aunt intentional.<sup>4</sup> Knobe argued that skill neglect occurs here because the act is morally *significant*. Thus, we designed a vignette in which not only difficulty could be held constant and skill could be manipulated but one in which the direction of moral valence — praiseworthiness or blameworthiness — could be varied as well. Knobe predicts skill neglect in both valence conditions (i.e., high intentionality ratings for unskilled actions that are either negative or positive). Models that postulate biased causality or intentionality judgments for blame judgments (e.g., Alicke, 2000) might also predict skill neglect, but only for the negative valence condition. Our analysis predicts no skill neglect in either valence condition. Once difficulty is held constant and no basic action stands in for the broader action of interest, people should follow their folk concept of intentionality, which treats skill as one necessary condition of intentional action regardless of moral valence.

## Study 6

### *Method*

With the above goals in mind, we designed a set of vignettes in which morally praiseworthy and morally blameworthy motives played out in the same action, thus holding difficulty constant. We created a context in which the transgression was somewhat more natural and representative of everyday interaction, which more often contains meanness and inconsiderateness than extreme transgressions such as murder. In addition, we manipulated the actor's level of skill. Thus, we administered four vignettes that constitute a 2 (moral valence: praiseworthy vs. blameworthy) × 2 (skill: some vs. none) design.

Participants were 120 undergraduate students who completed a larger computer-presented questionnaire and received partial course credit in return. Each received one of four vignettes.

The background story common to all conditions was that George and his sister Lena reunite at their parents' house for Thanksgiving. Whereas George just got into medical school, Lena is unhappy in her marriage and recently lost her job. Over the course of the day, George and Lena get into a number of heated arguments. Later in the afternoon they play a game of darts. They split the first two games, and the third game is close until the end. “Who will win

---

<sup>4</sup> The saving lives vignette is afflicted with the same problems as the shooting aunt vignette — the prominence of pulling the trigger and the verb choice of “hitting the device” (rather than, for example, “hitting the power switch of the device”). The important point is, however, that people treat the *morally* praiseworthy saving lives vignette just like the *morally* blameworthy killing vignette, not like the (simply praiseworthy) achievement vignette.

comes down to George's last shot. If he hits 11 or more, he wins; if he hits less than 11, Lena wins." An image of a dart board was included in the vignette.<sup>5</sup>

We manipulated moral valence by varying George's goal and its outcome. In the praiseworthy condition, "George thinks of the difficult time Lena is having, and he really wants to let her win." After his shot, "the dart lands in the 6-point segment. Lena wins the game and is very happy." In the blameworthy condition, "George doesn't care that Lena is having a difficult time; he really wants to beat her." After his shot, "the dart lands in the 14-point segment. George triumphs in his victory while Lena is sad."

The actor's skill was manipulated as follows. In the no-skill condition, "George has little control over his shots. He tries to aim at a region of the darts board in which he would probably make [fewer/more] than 11 points. As he sets up his shot, however, he loses his balance, the dart slips out of his hand and wobbles toward the board, and... the dart lands in the [6-point/14-point] segment." In the skill condition "He aims the dart at a region of the board in which he would probably make [fewer/more] than 11 points. He sets up his shot and... the dart lands in the [6-point/14-point] segment."

Participants first answered the intentionality question, "Did George hit the [6-point/14-point] segment intentionally?", on a 4-point scale from 0 (Certainly no) to 3 (Certainly yes), with the numbers in between labeled Probably no and Probably yes. The scale was presented as continuous to make the judgment somewhat easier for participants, but for analysis purposes, the two no answers and the two yes answers were combined. (A continuous analysis leads to identical results.) Participants also answered an evaluation question: "Do you approve or disapprove of George's action?" on a 4-point scale from -2 (Very much disapprove), -1 (Somewhat disapprove), +1 (Somewhat approve), to +2 (Very much approve).

Finally, we devised an intuitively meaningful manipulation check for the skill factor that also served as an equality check for the action's difficulty. An agent's skill and an action's difficulty jointly result in the agent's likelihood of performing the action. Because the vignette required that the action was performed, we asked participants to judge how likely it is that George would perform it again. "Their cousin Frank feels that George's shot was lucky. He challenges George to hit the [14-point/6-point] segment again. How likely is it that George will hit the [14/6] again?" The response scale ranged from 0 (Extremely unlikely) to 5 (Extremely likely), with the numbers in between labeled Very Unlikely, Somewhat Unlikely, Somewhat Likely, and Very Likely. Manipulation of the skill factor should result in reliable differences on this likelihood rating whereas likelihood ratings should not differ across action valence.

## Results

The valence manipulation check showed that people approved of the benevolent action more ( $M = 1.2$ ) than of the mean action ( $M = 0.3$ ),  $F(1,114) = 25.5, p < .001$ . By contrast, valence did not have an impact on the agent's likelihood of repeating the same shot ( $F = 1, d = 0.18$ ), suggesting that people viewed the praiseworthy and the blameworthy act as equally

---

<sup>5</sup> The cutoff of 11 was chosen because it splits the 20 possible segments in two halves and makes the over/under appear equally difficult. Strictly speaking, in many games of darts the 11 would not be the midpoint of the *points* one can achieve (because double and triple scores increase the possible points all the way to 60), but we refrained from explaining this complication to the participants. We ran a variant of this study in which the over/under cutoff was 20 points, which is close to the median of the entire possible point range, and we found the same results.

difficult. The skill manipulation had an expected effect on the likelihood rating, with the unskilled agent being less likely to repeat the same shot ( $M = 1.5$ ) than the skilled agent ( $M = 2.2$ ),  $F(1, 116) = 16.5$ ,  $d = 0.73$ .

As predicted, intentionality judgments were sensitive to the skill manipulation: People more often judged the action intentional when the agent had skill (85%) than when he did not (27%). Most important, the effect of skill on intentionality judgments was as strong for the blameworthy action (87% vs. 23%,  $d = .84$ ) as for the praiseworthy action (83% vs. 30%,  $d = .61$ ). In a loglinear analysis predicting intentionality judgments from skill and valence, the skill effect was substantial and significant,  $z = 5.9$ ,  $p < .001$ ,  $d = 0.76$ , whereas the valence effect was not,  $z = 0.1$ ,  $d = 0.01$ , nor was the valence  $\times$  skill interaction,  $z = 0.63$ ,  $d = 0.08$ .

### *Discussion*

The results showed that people distinguish between low and high skill independent of the behavior's valence — they are just as sensitive to the lack of skill for positive behaviors as they are for negative behaviors (and this sensitivity is strong in both cases). However, the degree of disapproval elicited by the negative valence condition was not overwhelming, as only 41% of participants in the blameworthy condition gave George's action a negative rating. It is therefore possible that the negative act was not blameworthy enough to show any skill neglect. However, when we repeated all analyses on only those participants in the blameworthy condition who expressed explicit disapproval ( $N = 24$  out of 60), the results were just as strong (77% saw the blameworthy action as intentional when performed with skill, 9% saw it as intentional when performed without skill).

Knobe's (2003b) hypothesis of skill neglect in morally significant judgments predicted high intentionality judgments, even in the absence of skill, for both the praiseworthy and the blameworthy action. In Study 6, neither action showed such skill neglect. The study design assured that there was no basic action performed intentionally (such as pulling the trigger) that could count as the broader action, and the two actions compared were equated for difficulty. Under these circumstances, no undue influence of moral evaluation on intentionality judgments remained.

### General Discussion

Knobe (2003b) challenged traditional theories of moral judgment and intentionality by suggesting that people judge morally significant behaviors as intentional even if a component of intentionality (such as skill) is missing — quite in contrast to neutral behaviors, which are judged unintentional when missing a component. The valence (and especially blameworthiness) of a behavior, he proposed, guides perceptions of intentionality, not the other way around, as was traditionally assumed. He provided evidence for this hypothesis by showing that intentionality judgments for an unskilled act of killing (or saving lives) are frequent whereas intentionality judgments for an unskilled act of hitting the bull's eye are infrequent. Similar intentionality discrepancies between valenced and neutral actions were reported by Nadelhoffer (2004, 2005).

In our studies we have identified two alternative processes that account for Knobe's findings. The first process is that broader actions (such as killing) can be accomplished by basic actions (such as pulling the trigger), and performing a basic action intentionally can count as performing the broader action intentionally (Studies 1 and 2). The second process is that people see killing as easier than hitting a bull's eye, and because easier actions require less skill to be performed intentionally, people consider killing intentional even when the actor displays low

skill (Studies 3 to 5). When difficulty is held constant, any difference in intentionality judgments disappears, both for morally blameworthy and morally praiseworthy actions (Studies 5 and 6).

These two processes can be integrated within a framework of an action's *scope*, which refers to the ease and number of ways in which the action can be accomplished (cf. Goldman, 1970; Wegner & Vallacher, 1986). An action with wide scope (e.g., work, talk) has two features: First, it can be accomplished by a variety of basic actions that count as the performance of that wide-scope action. Second, among the variety of basic actions that constitute an action with wide scope, several basic actions will be relatively easy, so the wide-scope action becomes easier as well and can be performed intentionally even with a low level of skill. By contrast, an action with narrow scope (e.g., simmer the sauce, break the 100-meter record) can be accomplished only by very few basic actions and requires considerable skill to do so. In each of Knobe's (2003b) studies, a morally significant action (killing, saving lives) was compared with a morally neutral action (hitting the bull's eye), but both morally significant actions had wide scope whereas the neutral action had narrow scope. As a result, killing and saving lives could be accomplished by a relatively easy basic action (pressing the trigger), which allowed the agent to perform these actions intentionally even at low levels of skill. Hitting the bull's-eye, by contrast, was not sufficiently constituted by pulling the trigger and was seen as more difficult, thus preventing the agent with low levels of skill to perform the action intentionally. The gunman did not have enough skill to intentionally hit the bull's eye, but he had enough skill to intentionally kill.

The present studies thus disconfirm the hypothesis that moral evaluations bias intentionality judgments. But beyond that, they teach us several lessons about both intentionality judgments and moral judgment.

### *Lessons Learned*

*Skill and difficulty.* Our results document that people are not only sensitive to levels of skill when judging intentionality but integrate the agent's skill with the action's difficulty. The amount of skill necessary for an action to be intentional is not some absolute amount but is lower for easy actions and higher for more difficult actions. This interplay between skill and difficulty was first described by Heider (1958) as the concept of *can*, which he defined as follows:

$$can = f(\text{person's ability} - \text{difficulty of environmental factors})$$

We may thus sharpen Malle and Knobe's (1997) model of the folk concept of intentionality by replacing the *skill* component with the *can* component. What is needed for a behavior to be seen as intentional is that the agent *can* perform the behavior (in a reliable and controlled manner), so she has to have enough skill to overcome the action's difficulty.

*Blame and intentionality.* Judgments of blame are not as strongly tied to intentionality judgments as was previously expected. In our studies, correlations between the two judgments ranged from  $r = .07$  to  $r = .39$ , with an average of  $r = .22$ . Admittedly, people typically examined actions of extreme valence — (murdering or injuring another person), which restrict the range of blame judgments. However, Study 6 examined actions of more moderate valence and showed that intentionality judgments were sensitive to variations in skill whereas blame judgments were sensitive to moral valence (but not vice versa), leaving only a correlation of  $r = .19$  between blame and intentionality. Moreover, in Study 3 blame and intentionality diverged when pitted against each other. We compared two actions, one more difficult and more blameworthy (shooting a person in the head) than the other (shooting the person in the arm). If intentionality followed blame, the more blameworthy action should show increased intentionality; if intentionality followed difficulty, the more blameworthy (but more difficult)

action should show *decreased* intentionality. The latter was the case. Of course, when we generally compare blame for behaviors that are clearly intentional with blame for behaviors that are clearly unintentional, the two judgments will be highly correlated. However, the present studies confirm that such a correlation holds because intentionality guides blame, not the other way around.

*Levels of action.* People simultaneously assess intentionality for basic and for broader actions and integrate the two as a function of the action's scope. In the case of wide-scope actions, the intentionality of a suitable basic action transfers to the broader action, whereas in the case of narrow-scope actions, such suitable basic actions are hard to find.

The notion that people respond differentially to behaviors with narrow or wide scope finds kinship with two related literatures. Action identification theory (e.g., Vallacher & Wegner, 1987) posits that any given action can be identified in multiple ways, from its concrete performance features to its abstract meaning, and that the identification level of an action can influence people's attention to and interpretation of it. Similarly, the linguistic category model (Semin, 2008) suggests that speakers use more or less abstract verbs to influence others' impressions of an individual or group. Descriptive action verbs such as "stare" or "push" denote specific actions of shorter temporal duration that are easily verifiable, whereas interpretive action verbs such as "deride" or "harm" denote broader classes of actions that are more difficult to verify, more open to interpretation (Semin & Fiedler, 1988). Consistent with these perspectives on behavior identification and description, our results illustrate how action descriptions at different levels of specificity influence judgments of intentionality. Thus, narrow-scope descriptions use descriptive action verbs (e.g., hitting the bull's eye) to refer to an action's concrete performance that is temporally circumscribed and easily verified, and as a result, criteria of intentionality apply precisely. Wide-scope descriptions use interpretive action verbs (e.g., kill) to refer to the action's broader meaning, which is harder to verify and leaves more room for interpretation of properties such as intentionality.

*Causal deviance.* The present findings help clarify the role of causal deviance in moral judgment (e.g., Chisholm, 1966). Previous studies have shown that judgments of both blame and praise are weakened when an agent's behavior is causally deviant — when the agent's intended outcome unfolds in a way that was not intended (Pizarro, Uhlmann, & Bloom, 2003). Our Studies 1 and 2 demonstrated that causal deviance sometimes leaves both intentionality and blame untouched, namely, when the deviance occurs *after* intentionally performing a basic action that counts as the broader action at issue. However, if causal deviance interferes with the basic action and indicates a lack of control, the broader action is less likely to be deemed intentional, even if it results in a severely negative outcome (see Mele & Cushman, 2008 for similar findings). Moreover, only if people view the causally deviant action as unintentional will they reduce assignments of blame.

*Jury decision making.* Knobe's (2003a, b) challenges to the standard intentionality model threaten the good sense of asking people to make *mens rea* judgments in legal proceedings (Nadelhoffer, 2006b). We should be wary of jurors' assessments of a defendant's intent if the negative valence of the presumed criminal act biases them toward "seeing" such intent. However, as the present findings meet the skill challenge and those reported elsewhere meet the side-effect challenge (Guglielmo & Malle, under review), we should maintain our confidence in jurors' ability to handle judgments of *mens rea* (Malle, 2006; Malle & Nelson, 2003). Moreover, our studies highlight the conditions under which such judgments can go awry, since, sadly, some of the methodological problems in Knobe's and other researchers' studies are mirrored in

procedural problems in the courtroom. When cases are presented with vague, wide-scope action verbs, generally unfavorable background information about the defendant, and limited judgment categories (e.g., intentionality—yes or no?), jurors will not be at their best in assessing the intentionality of the defendant's action. Conversely, such assessments will be more reliable with precise formulations of the actions in question, careful selection of what background is relevant (i.e., information that helps infer the mental states at hand, not general character), and permission to use multiple conceptual categories to describe the action in question (e.g., willingly, knowingly, with luck).

### *Limitations and Future Directions*

One limitation of our studies is the reliance on especially negative behaviors, usually murder or severe injury. Given that such scenarios were the source of the skill challenge (Knobe, 2003b; Nadelhoffer, 2004, 2005), we had to make them our primary focus in order to fully account for the phenomenon. However, Study 6 hints at the worthwhile aim for future research to explore more fully the relationship between difficulty, skill, intentionality, and blame/praise across a range of morally negative and positive behavior.

Second, whereas our findings counter the claim that judgments of moral blame (or praise) guide intentionality judgments, we did not assess the actual timing of these judgments. For highly negative behaviors, blame and intentionality will often co-occur, and disentangling their temporal order and causal direction can be difficult. One research strategy for such disentangling is to examine people's reaction times of making the various judgments in response to an observed behavior. We are currently designing and running studies that use such a reaction time paradigm and offer promise for further understanding the interplay between blame and intentionality (Guglielmo & Malle, 2008).

The literature will also benefit from further study of people's spontaneous frames and descriptions of negative and positive behaviors. As we have shown, intentionality judgments are highly sensitive to the precise way in which a given action is described. But how do people themselves describe these actions? A negative behavior is less likely to be seen as intentional when framed as having narrow (as opposed to broad) scope. Speakers may therefore opt to describe such behaviors as narrow or broad depending on whether they want to convince their audience of the action's intentionality or unintentionality. We suspect that, in the courtroom, attorneys for the prosecution will favor broad descriptions to heighten the impression of intentionality, whereas attorneys of the defense will favor narrow descriptions to weaken the impression of intentionality. Previous studies (Knobe, 2003b; Nadelhoffer, 2004, 2005) mimicked the prosecution. We hope that our studies have given justice to the other side and have demonstrated the subtlety with which ordinary people think about intentionality and blame.

## References

- Adams, F. (1986). Intention and intentional action: The simple view. *Mind and Language*, 1, 281-301.
- Alicke, M. D. (2000). Culpable control and the psychology of blame. *Psychological Bulletin*, 126, 556-574.
- Chisholm, R. (1966). Freedom and action. In K. Lehrer (Ed.), *Freedom and Determinism*. New York: Random House.
- Cushman, F. & Mele, A. R. (2008). Intentional action: two-and-a-half folk concepts? In J. Knobe and S. Nichols (Eds.), *Experimental Philosophy*. New York: Oxford University Press.
- Goldman, A. (1970). *A theory of human action*. Englewood Cliffs, NJ: Prentice-Hall.
- Guglielmo, S. & Malle, B. F. (2008). *The priority of intentionality or blame: Evidence from a reaction time experiment*. Poster presented at the meeting of the Society for Personality and Social Psychology, Albuquerque, NM.
- Guglielmo, S. & Malle, B. F. (under review). *Can unintended side effects be intentional? Solving a puzzle in people's judgments of intentionality and morality*.
- Hamilton, V. L. (1978). Who is responsible? Towards a social psychology of responsibility attribution. *Social Psychology*, 41, 316-328.
- Heider, F. (1958). *The psychology of interpersonal relations*. New York: Wiley.
- Knobe, J. (2003a). Intentional action and side effects in ordinary language. *Analysis*, 63, 190-193.
- Knobe, J. (2003b). Intentional action in folk psychology: An experimental investigation. *Philosophical Psychology*, 16, 309-324.
- Malle, B. F. (2006). Intentionality, morality, and their relationship in human judgment. *Journal of Cognition and Culture*, 6, 61-86.
- Malle, B. F. & Knobe, J. (1997). The folk concept of intentionality. *Journal of Experimental Social Psychology*, 33, 101-121.
- Malle, B. F., Moses, L. J., & Baldwin, D. A. (2001). *Intentions and intentionality: Foundations of social cognition*. Cambridge, MA: MIT Press.
- Malle, B. F. & Nelson, S. E. (2003). Judging mens rea: The tension between folk concepts and legal concepts of intentionality. *Behavioral Sciences and the Law*, 21, 563-580.
- McCann, H. J. (2005). Intentional action and intending: Recent empirical studies. *Philosophical Psychology*, 18, 737-748.
- Mele, A. R. (1992). *Springs of action: Understanding intentional behavior*. New York: Oxford University Press.
- Mele, A. R. (2003). Intentional action: Controversies, data, and core hypotheses. *Philosophical Psychology*, 16, 325-340.
- Mele, A. R. & Moser, P. K. (1994). Intentional action. *Noûs*, 28, 39-68.
- Mele, A. R. & Sverdlik, S. (1996). Intention, intentional action, and moral responsibility. *Philosophical Studies*, 82, 265-287.
- Nadelhoffer, T. (2004). The Butler problem revisited. *Analysis*, 64, 277-284.
- Nadelhoffer, T. (2005). Skill, luck, control, and intentional action. *Philosophical Psychology*, 18, 341-352.

- Nadelhoffer, T. (2006a). Desire, foresight, intentions, and intentional actions: Probing folk intuitions. *Journal of Cognition and Culture*, 6, 133-157.
- Nadelhoffer, T. (2006b). Bad acts, blameworthy agents, and intentional actions: Some problems for jury impartiality. *Philosophical Explorations*, 9, 203-220.
- Nichols, S. & Mallon, R. (2006). Moral rules and moral dilemmas. *Cognition*, 100, 530-542.
- Nichols, S. & Ulatowski, J. (2007). Intuitions and individual differences: The Knobe effect revisited. *Mind and Language*, 22(4), 346-365.
- Ohtsubo, Y. (2007). Perceiver intentionality intensifies blameworthiness of negative behaviors: Blame-praise asymmetry in intensification effect. *Japanese Psychological Research*, 49, 100-110.
- Pizarro, D. A., Uhlmann, E., & Bloom, P. (2003). Causal deviance and the attribution of moral responsibility. *Journal of Experimental Social Psychology*, 39, 653-660.
- Semin, G. R. (2008). Language puzzles: A prospective retrospective on the Linguistic Category Model. *Journal of Language and Social Psychology*, 27, 197-209.
- Semin, G. R. & Fiedler, K. (1988). The cognitive functions of linguistic categories in describing persons: Social cognition and language. *Journal of Personality and Social Psychology*, 54, 558-568.
- Sousa, P. & Holbrook, C. (under review). Folk concepts of intentional action (in the contexts of amoral and immoral luck).
- Shaver, K. G. (1985). *The attribution of blame: Causality, responsibility, and blameworthiness*. New York: Springer.
- Tedeschi, J. T., & Reiss, M. (1981). Verbal strategies as impression management. In C. Antaki (Ed.), *The psychology of ordinary social behaviour* (pp. 271-309). London: Academic Press.
- Thompson, S. C., Armstrong, W., & Thomas, C. (1998). Illusions of control, underestimations, and accuracy: A control heuristic explanation. *Psychological Bulletin*, 123, 143-161.
- Vallacher, R. R. & Wegner, D. M. (1987). What do people think they're doing? Action identification and human behavior. *Psychological Review*, 94, 3-15.
- Wegner, D. M. & Vallacher, R. R. (1986). Action identification. In R. M. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (pp. 550-582). New York: Guilford.
- Weiner, B. (1986). *An attributional theory of motivation and emotion*. New York: Springer.
- Weiner, B. (1995). *Judgments of responsibility: A foundation for a theory of social conduct*. New York: Guilford.