

A Knobe Effect for Belief Ascriptions*

James R. Beebe (University at Buffalo)

The Review of Philosophy and Psychology (forthcoming)

Knobe (2003a, 2003b, 2004b) found that people are more likely to attribute intentionality to agents whose actions resulted in negative side-effects than to agents whose actions resulted in positive ones. Subsequent investigation has extended this result to a variety of other folk psychological attributions. The present article reports experimental findings that demonstrate an analogous effect for belief ascriptions. Participants were found to be more likely to ascribe belief, higher degrees of belief, higher degrees of rational belief, and dispositional belief to agents in central Knobe effect cases who bring about negative side-effects than to agents who bring about positive ones. These findings present a significant challenge to widely held views about the Knobe effect, since many explanations of it assume that agents in contrasting pairs of Knobe effect cases do not differ with respect to their beliefs. Participants were also found to be more confident that knowledge should be attributed than they were that belief or dispositional belief should be attributed. This finding strengthens the challenge that Myers-Schulz and Schwitzgebel (forthcoming) have launched against the traditional view that knowledge entails belief.

keywords: Knobe effect, experimental philosophy, belief ascription, folk psychology

* Thanks to Maria Capolupo, Sean Carey, Mattias Carosella, Phillip Collins, Danielle Curtin, Rachel Pazda, Jordan Piridy, and Paul Poenicke who served as research assistants on this project. Thanks also to Mark Alfano, two anonymous reviewers at *The Review of Philosophy and Psychology*, and audience members at Eindhoven University of Technology and the 2012 meeting of the Society for Exact Philosophy for helpful comments on previous drafts of this paper.

1. Introduction

In what has become known as the ‘Knobe effect’ (or ‘side-effect effect’) Joshua Knobe (2003a, 2003b, 2004b) discovered that people are more likely to say that an agent intentionally brought about a result that was the side-effect of the agent’s primary action if the side-effect was negative than if it was positive. Additional studies have shown that this effect is not confined to attributions of intentional action but is in fact a pervasive feature of folk psychological attributions. Pettit and Knobe (2009), for example, found that participants are more likely to say that an agent *decided to* bring about a side-effect, *was in favor of* or *opposed to* a side-effect, or *advocated* a side-effect if the side-effect was negative than if it was positive. Tannenbaum, Ditto and Pizarro (2007) found a similar effect with *desire*. More recently, Beebe and Buckwalter (2010) and Beebe and Jensen (2012) have discovered that participants are more likely to think an agent *knew* a side-effect was going to occur if the side-effect was negative than if it was positive. The present article reports the results of five experiments that extend this line of inquiry into the realm of belief ascriptions.¹

Knobe (2003a, p. 191) famously began his investigation of the Knobe effect by presenting participants with either the help or the harm version of the following vignette:

ENVIRONMENT: 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, *and/but* it will also *help/harm* the environment.” The chairman of the board answered, “I don’t care at all about *helping/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 *helped/harmed*.

¹ Additional studies (e.g., Cova and Naar forthcoming) have uncovered a similar effect with means rather than side-effects.

Participants were asked whether the chairman intentionally helped or harmed the environment. Knobe found that 82% of participants in the harm condition answered that the chairman intentionally harmed the environment, but only 23% of those in the help condition thought he intentionally helped it. Knobe (2003a, p.190) began the article that introduced these results with the following words:

The chairman of the board of a company has decided to implement a new program. He believes

(1) that the program will make a lot of money for his company

and

(2) that the program will also produce some other effect *x*.

In other words, Knobe assumed that the wording of ENVIRONMENT made it obvious that the chairman believed the side-effect would occur in both the help and the harm conditions. Practically every philosopher and psychologist working on the Knobe effect has followed suit and has assumed that the chairman both knew and believed that the side-effects in question would occur.² As we will see in Section 7 below, many of the explanations offered of the Knobe effect depend upon this assumption in important ways. However, none of these scholars has any empirical basis for assuming the chairman's belief states are the same in both conditions. What is more, the results reported below show this assumption to be false.

Sections 2, 3, 4, and 6 report the results of four experiments in which participants were asked (a) whether the central protagonists in contrasting Knobe effect cases believed that certain side-effects would occur, (b) how strongly the protagonists believed the side-effects would occur, (c) how strongly the agents should believe they would occur, and (d) whether the

² Cf., e.g., Hindriks (2008), Knobe (2004a, 2006), Knobe and Burra (2006), Leslie, Knobe and Cohen (2006), Doris, Knobe and Woolfolk (2007), Knobe and Doris (2010), Nadelhoffer (2004a, 2004b, 2006), Machery (2008), McCann (2005), Mele (2006), Mele and Cushman (2007), Cushman and Mele (2007), Pettit and Knobe (2009), Sverdlik (2004).

protagonists were disposed to believe that the side-effects would occur. When presented with Knobe's original ENVIRONMENT vignette, participants were more inclined to attribute belief, greater strength of belief, greater strength of rational belief, and dispositional belief to protagonists whose actions brought about negative side-effects than to those who brought about positive side-effects. Similar effects were found in some but not all of the additional vignettes employed, all of which were drawn from previously published studies of the Knobe effect. After noticing that participants were somewhat less inclined to attribute belief in cases where other researchers had found them willing to attribute knowledge, two additional experiments (described in Sections 5 and 6) were constructed that asked participants about both the knowledge and beliefs of the protagonists. In ENVIRONMENT and one other case participants were found to be significantly more confident that knowledge should be attributed than that belief should be attributed. These results bolster Myers-Schulz and Schwitzgebel's (forthcoming) contention that knowledge may not entail belief, as has been traditionally assumed.

2. Experiment 1

2.1. Methods

2.1.1. Participants

400 undergraduate students (48% female, predominantly Caucasian, average age = 20) from a large, public university in the northeastern United States participated in Experiment 1 on a voluntary basis. They were neither paid nor given course credit for their participation.

2.1.2. Materials and procedure

In a between subjects design, 145 participants were presented with either the help or the harm

version of Knobe's (2003a) original ENVIRONMENT vignette. They were asked to indicate on a seven-point Likert scale (ranging from -3 to 3) the extent to which they agreed or disagreed with the following claim: "The chairman believed that the new program would help/harm the environment." '-3' was labeled 'Strongly Disagree,' '0' 'Neutral,' and '3' 'Strongly Agree.'

Following Beebe and Jensen (2012), an additional 84 participants were given either the better or the worse version of the following, modified version of Knobe's (2004a) vignette:

MOVIES: The Vice-President of a movie studio was talking with the CEO. The Vice-President said: "We are thinking of implementing a new policy. If we implement the policy, it will increase profits for our corporation, but it will also make our movies *better/worse* from an artistic standpoint." The CEO said: "Look, I don't care one bit about making our movies *better/worse* from an artistic standpoint. All I care about is making as much profit as I can. Let's implement the new policy." They implemented the policy. Sure enough, the policy made the movies *better/worse* from an artistic standpoint.³

In contrast to ENVIRONMENT, the relevant sort of good and bad in MOVIES is non-moral and aesthetic. Participants were asked to indicate on a seven-point Likert scale the extent to which they agreed or disagreed with the following claim: "The CEO believed that the new policy would make the movies *better/worse* from an artistic standpoint." As in the previous case, '-3' was labeled 'Strongly Disagree,' '0' 'Neutral,' and '3' 'Strongly Agree.'

Knobe and Mendlow (2004, p. 24) constructed a vignette in which an agent brings about a negative side-effect but is not blameworthy for doing so. Following Beebe and Jensen (2012), 87 participants were presented with either the increase or the decrease version of the following, modified version of Knobe and Mendlow's (2004) vignette:

³ Italics did not appear in the research materials and are used here only for purposes of explanation.

SALES: Susan is the president of a major computer corporation. One day, her assistant comes to her and says, “We are thinking of implementing a new corporate restructuring plan. It will simplify our corporate structure, *and/but* it will also *increase/decrease* sales in New Jersey for the next quarter.” Susan says “I don’t care about what happens in the next quarter. We need to simplify our corporate structure. Let’s implement the new plan.” They implemented the new plan. Sure enough, sales in New Jersey *increased/decreased* in the following quarter.

Participants were asked to indicate on a seven-point Likert scale the extent to which they agreed or disagreed with the claim “Susan believed that the new plan would *increase/decrease* sales in New Jersey in the next quarter.” The scale was labeled as before.

In an effort to show that nonconscious moral judgments can affect attributions of intentionality, Knobe (2007) constructed a pair of cases that were designed to elicit a conflict between participants’ immediate, intuitive judgments and their more reflective, conscious judgments. Following Beebe and Jensen (2012), 84 participants were given either the fulfill or the violate version of the following, modified vignette from Knobe (2007):

NAZI: In Nazi Germany, there was a law called the “racial identification law.” The purpose of the law was to help identify people of certain races so that they could be rounded up and sent to concentration camps. Shortly after this law was passed, the CEO of a small corporation decided to make certain organizational changes. The vice-president of the corporation said: “By making those changes, you’ll definitely be increasing our profits. But you’ll also be *violating/fulfilling* the requirements of the racial identification law.” The CEO said: “I don’t care one bit about that. All I care about is making as much

profit as I can. Let's make those organizational changes!" As soon as the CEO gave this order, the corporation began making the organizational changes.

Participants were asked to indicate the extent to which they agreed or disagreed with the claim "The CEO believed that the organizational changes would *violate/ fulfill* the requirements of the law," and their responses again were registered on a 7-point scale ranging from -3 to 3.

2.2. Results

Significant differences between participant responses in pairs of contrasting conditions were found in three of the four cases (cf. Figure 1 and Table 1). Participants presented with ENVIRONMENT were significantly more likely to think the chairman believed the side-effect would come about when the environment was being harmed than when it was being helped. Participants were twice as likely to answer 'Strongly Agree' (i.e., '3') to the question of whether the chairman believed the environment would be harmed (40%) than they were to give the same answer in response to the question of whether he believed it would be helped (19%).




Figure 1. Mean participant responses in each of the eight conditions of Experiment 1. Error bars in all graphs represent 95% confidence intervals. In every graph an ‘*’, ‘**’ or ‘***’ by itself indicates that the mean differs significantly from the neutral midpoint at either the .05, .01 or the .001 level. An ‘*’, ‘**’ or ‘***’ with a bracket indicates a statistically significant difference between pairs of conditions at either the .05, .01 or the .001 level.⁴

	ENVIRONMENT		MOVIES		SALES		NAZI	
	Help	Harm	Better	Worse	Increase	Decrease	Fulfill	Violate
Mean	.27	1.27	-.85	.39	.47	.43	.37	1.28
SD	2.00	1.84	1.76	2.00	1.83	2.12	2.06	1.89

Table 1. Mean participant responses and standard deviations for each of the four pairs of conditions in Experiment 1.

⁴ Environment: Mann-Whitney $U = 1843, p < .01, r = -.25$. Movies: $U = 568.5, p < .01, r = -.31$. Sales: $U = 919.5, p > .05, r = -.02$. Nazi: $U = 651, p < .05, r = -.23$. All significance tests two-tailed. Pairwise tests of statistical significance were used instead of a two-way ANOVA because the differences between ENVIRONMENT, MOVIES, SALES, and NAZI render them separate pairs of cases more than different levels of a common independent variable.

In a similar fashion, participants in the worse condition of MOVIES were more likely to think the CEO believed that the new policy would make the movies worse than were participants in the better condition to think he believed it would make them better. 52% of participants in the worse condition answered '1,' '2' or '3,' compared to 22.5% in the better condition. 30% of participants in the worse condition chose a negative answer of '-1,' '-2' or '-3,' while 60% did so in the better condition. In contrast to the results from ENVIRONMENT and MOVIES, participants in the decrease condition of SALES were not more likely to think the central protagonist believed the new plan would decrease sales in New Jersey in the next quarter than were participants in the increase condition to think the protagonist believed it would increase them. However, participants in both conditions were significantly more likely to give a positive answer (increase = 56%, decrease = 55%) than a negative (23%, 32%) one.

Participants in the violate condition of NAZI were more likely to think the CEO believed organizational changes would violate the requirements of the racial identification law than participants in the fulfill condition were to think he believed they would fulfill the law. Participants in both conditions were more likely to give positive answers (fulfill = 49%, decrease = 72%) than negative answers (37%, 16%), but there was a higher percentage of positive answers in the violate condition than in the fulfill condition. Effect sizes in the three cases where statistically significant results were obtained ranged from small to medium, which is comparable to other findings in the Knobe effect literature.⁵ No gender effects were observed.

⁵ The apparent size of a given Knobe effect has been shown to depend in part upon the kind of prompt question participants are asked. For example, Beebe and Jensen (forthcoming, secs. 1-2) found small effect sizes when using 7-point Likert scales but large effect sizes when using forced-choice formats, even though other features of the experimental materials were held constant.

2.3. Discussion

Readers should beware of falling prey to the ‘less is better effect’ (Hsee 1998) when evaluating the significance of foregoing results. If I had given participants only Knobe’s original ENVIRONMENT case—or if I had followed common practice in experimental philosophy and reported only the results that were most beneficial to my case—the experimental philosophy community would have interpreted me as conclusively showing that there is a Knobe effect for belief ascriptions. However, since I used (and reported the results from) four cases, and one of them did not give rise to the predicted effect, some readers may be tempted to conclude that I have not established my case as strongly as Knobe did. After all, my record of success is only 75%, whereas Knobe’s was 100%. However, just as adding a few broken dishes to a set of flawless dinnerware should not reduce the value of the flawless dinnerware (Hsee 1998), using additional vignettes in which the target phenomenon was not found should not detract from the significance of the fact that it was found in the paradigmatic ENVIRONMENT case. In fact, I would argue that the fact that a Knobe effect was not found for SALES is most likely to be an indication that it is not as good an example of a Knobe effect case as ENVIRONMENT rather than being an indication that the data offered in favor of a Knobe effect for belief ascriptions are mixed.

With these qualifications in mind, I think we can conclude that the results from Experiment 1 show that Knobe (2003a) and others were too hasty in assuming that the chairman (and other similar protagonists) would naturally be interpreted as believing that the primary action in question would produce the side effect action, regardless of which condition he was in. As will be explained in Section 7 below, these findings undermine many explanations of the

Knobe effect that crucially depend upon this assumption and provide the basis for a more adequate interpretation of the full range of Knobe effects.

The data obtained in Experiment 1 also replicate and extend the findings of Sverdlik (2004), who in two separate experiments presented participants with only the harm version of ENVIRONMENT and asked them whether the chairman believed he would harm the environment. In a forced choice format 62% of participants in the first experiment answered ‘Yes,’ and 38% answered ‘No.’ In the second experiment, 82% answered ‘Yes,’ and 18% answered ‘No.’ Experiment 1 shows that Sverdlik’s results are robust and can be found in cases beyond the basic ENVIRONMENT case.

3. Experiment 2

Experiment 1 treated belief in roughly a bivalent fashion—i.e., as something that either is or is not present—rather than as something that comes in degrees. Participants were asked to express their degree of agreement that a given belief was present, but strictly speaking participants were not able (or at least were not supposed) to ascribe a partial degree of belief to the protagonists in the vignettes. According to the intended interpretation of the research materials, a high rating of ‘3’ should indicate a high degree of confidence on the part of the participant that a belief was present, whereas a middling rating of ‘1’ should indicate a middling degree of confidence that this was the case. However, it is possible that some participants used the Likert scale to indicate their thoughts about the degree of strength of the protagonist’s belief itself. For example, a participant might select ‘1’ in order to express a very confident judgment that the protagonist had a rather weak degree of belief. In order to allow participants to express judgments about the degree of belief possessed by the protagonists in the experimental cases and to approach the

question of how normative assessments affect belief attributions from a different angle, participants in Experiment 2 were directed to think about protagonists' degrees of belief.

3.1. Methods

3.1.1. Participants

330 undergraduate college students (62% female, predominantly Caucasian, average age = 21) from a large, public university in the northeastern United States participated in this study on a voluntary basis. They were neither paid nor given course credit for their participation.

3.1.2. Materials and procedure

In a between subjects design each participant was given one of the vignettes employed in Experiment 1 and asked one of the following questions:

- How strongly do you think the chairman believed that the new program would *help/harm* the environment?
- How strongly do you think the CEO believed that the new policy would make the movies *better/worse*?
- How strongly do you think Susan believed that the new plan would *increase/decrease* sales in New Jersey for the next quarter?
- How strongly do you think the CEO believed that the organizational changes would *violate/fulfill* the requirements of the law?

Participants were asked to indicate their responses on a seven-point scale ranging from -3 to 3. In the chairman case '-3' was labeled 'He strongly believed that it would *not* help/harm the environment' (emphasis in test item), '0' was labeled 'He didn't believe one way or the other,'

and ‘3’ was labeled ‘He strongly believed that it would help/harm the environment.’ The scales used with the other vignettes were anchored in an analogous fashion.

3.2. Results

A statistically significant difference was found in three of the four cases, but only in ENVIRONMENT was the difference in the prediction direction (cf. Figure 2 and Table 2).⁶ In ENVIRONMENT, none of the participants in the help condition chose answers ‘-3,’ ‘-2’ or ‘-1’—i.e., answers in the direction of ‘He strongly believed that it would *not* help the environment’—while 91% of them chose ‘0,’ representing a lack of belief one way or the other. In the harm condition, by contrast, 83% of participants chose ‘1,’ ‘2’ or ‘3.’ The effect size in this case was quite large.




Figure 2. Mean participant responses in each of the eight conditions of Experiment 2.

⁶ Environment: Mann-Whitney $U = 221, p < .001, r = -.71$. Movies: $t(55) = 2.81, p < .001, r = .35$. Sales: $t(49) = 2.12, p < .05, r = .29$. Nazi: $t(75) = 1.07, p > .05, r = .12$. Independent samples t -tests were used for the latter three cases because the data distributions more closely approximate normal distributions.

	ENVIRONMENT		MOVIES		SALES		NAZI	
	Help	Harm	Better	Worse	Increase	Decrease	Fulfill	Violate
Mean	.16	1.93	.6	.07	.54	.14	1	.66
SD	.53	1.16	1.08	.51	1.14	.41	1.28	1.52

Table 2. Mean participant responses and standard deviations for each of the four pairs of conditions in Experiment 2.

A small but statistically significant difference between participant responses in the better and worse conditions of MOVIES was found, but not in the expected direction. Although participants in both conditions chose ‘He didn’t believe one way or the other’ more than any other answer, more participants chose ‘1,’ ‘2’ or ‘3’ in the better condition (45%) than in the worse condition (5%). This finding runs contrary to the results of Experiment 1, in which protagonists who had been shown by Beebe and Buckwalter (2010) and Beebe and Jensen (2012) to be more likely viewed as knowing they would produce certain side-effects were also viewed as believing that they would do so.

Participant responses in the increase and decrease conditions of SALES approximate those in MOVIES, with participants in the increase condition on average attributing a slightly higher degree of belief than participants in the decrease condition. Participants in the fulfill and violate conditions of NAZI did not differ in the strength of belief they attributed to the main protagonist, although there was a wider distribution of responses in both conditions than in the three previous cases and participants were in general more likely to give positive answers (fulfill = 65%, violate = 59%) than negative ones (6%, 20%). Gender differences approached but failed

to reach statistical significance, with male response means being slightly higher in both the positive and negative conditions.⁷

3.3. Discussion

Participant responses to the central ENVIRONMENT case confirmed antecedent predictions, thereby showing that there can be a Knobe effect for ascriptions of degree of belief. However, despite the fact that participants in Experiment 1 were more likely to attribute belief to the central protagonists in three of four Knobe effect cases, participants in Experiment 2 were not on the whole more likely to attribute stronger degrees of belief to those protagonists—even though the same cases were used in both experiments. The apparent inconsistency between these results is surprising. Almost every other study of the Knobe effect suggests that protagonists who bring about side-effects that are negative are more likely to be viewed as having brought those effects about intentionally and knowing they would do so. If a protagonist has been viewed as knowing that a given side-effect would occur, one would expect him to be viewed as believing that it would do so.

One important difference between the protagonists in ENVIRONMENT and the protagonists in some of the other cases is that the chairman does not seem at all justified in harming the environment in order to increase his company's profits, whereas Susan certainly seems justified in decreasing sales in New Jersey in the short run in order to simplify her corporation's internal structure. Presumably, the restructuring plan will lead to a long-term increase in profits. Indeed, when Phelan and Sarkissian (2008) asked participants to rate the degree of praise or blame Susan deserved for decreasing sales and to say whether they thought

⁷ Mean male response in positive conditions: .52. Mean female response in positive conditions: .24. Male negative mean: .92. Female negative mean: .38. $F(1, 215) = 3.418, p = .066, \text{partial } \eta^2 = .016.$

that decreasing sales was bad, 94.5% of participants indicated that Susan deserved neither praise nor blame, and only 14% thought that decreasing sales was bad. These data contrast sharply with the praise and blame ratings obtained by Knobe (2003a) for ENVIRONMENT. When Knobe asked participants to indicate on a scale from 0 to 6 how much praise or blame the chairman deserved for helping or harming the environment, the mean blame rating in the harm condition was 4.8, while the mean praise rating was 1.4. In a similar fashion, the CEO in MOVIES might well be justified in increasing his company's profits at the expense of the artistic quality of the movies it produces. A chief executive clearly has a duty to make his corporation profitable, but it is not clear that a movie industry executive has a duty to make movies that are good from an artistic perspective. In other words, the unexpected data obtained in these two cases might stem from the fact that these cases are not as similar to ENVIRONMENT as they might at first appear. The difference between the results obtained in Experiment 1 and those in Experiment 2 may also indicate that participants in Experiment 1 were understanding the probe questions in the intended manner—i.e., as asking participants about how confident they were that the protagonists believed the side-effect would occur rather than asking participants how confident the protagonists were that the side-effect would occur.

4. Experiment 3

In order to investigate belief ascriptions in Knobe effect cases from an additional angle, participants in Experiment 3 were given the same vignettes employed in the previous two experiments, but this time they were asked to indicate how strongly the protagonists should believe the side-effects in question would come about as a result of their primary actions. Results from Experiment 3 fell more in line with previous expectations.

4.1. Methods

4.1.1. Participants

334 undergraduate college students (61% female, predominantly Caucasian, average age = 20) from a large, public university in the northeastern United States participated in this study on a voluntary basis. They were neither paid nor given course credit for their participation.

4.1.2. Materials and procedure

In a between subjects design, participants were given one of the same vignettes used in Experiments 1 and 2, except that the following pieces of information were removed:

- They started the new program. Sure enough, the environment was *helped/harmed*.
- They implemented the policy. Sure enough, the policy made the movies *better/worse* from an artistic standpoint.
- They implemented the new plan. Sure enough, sales in New Jersey *increased/decreased* in the following quarter.
- As soon as the CEO gave this order, the corporation began making the organizational changes.

These bits of ‘hindsight’ information about how things turned out were omitted because participants were asked how strongly protagonists should believe the projected events would occur, and it would be infelicitous to ask participants this question if it was clear the protagonists already knew they did occur.⁸ Following Phelan (forthcoming), the following questions were used as proxies for an agent’s degree of rational belief:

⁸ Of course, the information deleted from NAZI was not hindsight information in the same sense as in the other cases, but it was deleted for the sake of consistency.

- How strongly do you think the chairman should believe that the new program will *help/harm* the environment?
- How strongly do you think the CEO should believe that the new policy will make the movies *better/worse*?
- How strongly do you think Susan should believe that the new plan will *increase/decrease* sales in New Jersey for the next quarter?
- How strongly do you think the CEO should believe that making the new changes will *fulfill/violate* the racial identification law?

Participants indicated their responses on a seven-point scale ranging from -3 to 3. In the chairman case ‘-3’ was labeled ‘He should strongly believe that it will *not* help/harm the environment’ (emphasis in test item); ‘0’ was labeled ‘He shouldn’t believe one way or the other’; and ‘3’ was labeled ‘He should strongly believe that it will help/harm the environment.’ The scales used with the other prompt questions were given analogous anchors.

4.2. Results

Significant differences were found in all four pairs of cases. Participants who were given the harm version of ENVIRONMENT were significantly more likely to say the chairman should believe the side-effect will occur than were participants who were given the help version (cf. Figure 3 and Table 3).⁹ The balance of participants who gave positive (‘1,’ ‘2’ or ‘3’) vs. negative (‘-1,’ ‘-2’ or ‘-3’) responses to ENVIRONMENT differed significantly between the help (19.5% vs. 61%) and harm (66% vs. 18%) conditions.

⁹ Environment: Mann-Whitney $U = 429.5$, $p < .001$, $r = -.46$. Movies: $U = 428.5$, $p < .001$, $r = -.42$. Sales: $U = 387$, $p < .001$, $r = -.5$. Nazi: $U = 462.5$, $p < .001$, $r = -.46$.




Figure 3. Mean participant responses in each of the eight conditions of Experiment 3.

	ENVIRONMENT		MOVIES		SALES		NAZI	
	Help	Harm	Better	Worse	Increase	Decrease	Fulfill	Violate
Mean	-0.76	0.89	0.92	0.13	-0.75	1.1	0.58	2.05
SD	1.68	1.53	1.53	0.73	1.7	1.51	1.78	1.83

Table 3. Mean participant responses and standard deviations for each of the four pairs of conditions in Experiment 3.

As in Experiment 2, participants who were given a version of MOVIES failed to respond to the prompt questions in the predicted manner. They were slightly more likely to say the CEO should strongly believe the new policy will make the movies better than they were to think he should strongly believe it will make them worse, even though the mean responses in both

conditions were near the neutral midpoint. Participants in the worse condition overwhelmingly favored the neutral response of '0,' with only 10% of them choosing any other answer.

Participant responses to the pair of SALES vignettes closely resemble those given in response to ENVIRONMENT, with only 23% of participants choosing a positive answer in the increase condition and 68% choosing one in the decrease condition. Participant data for the two NAZI cases are similar to data obtained in other investigations of the Knobe effect and the epistemic side-effect effect. The majority of participants in the violate condition (74%) chose '3' with relatively few participants chose any other answer. Participant responses in the fulfill condition were spread somewhat evenly around a common answer in the middle. Effect sizes in all four cases were medium to large. No gender differences were found.

4.3. Discussion

In three out of four cases in Experiment 3 participants were more inclined to think the central protagonists should believe the predicted side-effects would occur when those side-effects were negative than when they were positive, even though the epistemological grounds for the predictions in question were seemingly identical across the paired conditions. In each case an apparently trustworthy source provides a key decision maker with nearly identical testimony about the expected outcome. As in each of the previous experiments, participant responses to ENVIRONMENT came out exactly as predicted.

Because Beebe and Jensen (2012) found a Knobe effect for knowledge attributions using ENVIRONMENT, MOVIES, SALES, and NAZI, and yet because participants in Experiments 1 through 3 were not always as keen on attributing belief in such cases, I began to wonder whether Myers-Schulz and Schwitzgebel's (forthcoming) contention that belief is not a necessary

condition for knowledge might find some unexpected support from the results described above. I had anticipated uncovering results comparable to those of Beebe and Jensen because knowledge is generally thought to entail belief. However, as Myers-Schulz and Schwitzgebel suggest, this traditional philosophical view may fail to reflect folk epistemic practice. In order to investigate the matter further, I formulated the following two experiments, in which participants were prompted to think about attributing both knowledge and belief to the protagonists in ENVIRONMENT, MOVIES, SALES, and NAZI.

5. Experiment 4

Myers-Schulz and Schwitzgebel (forthcoming) presented participants with what they describe as five cases of “determinate, propositional knowledge that either are not cases of belief or are, at most, ‘in-between’ cases of belief in which the subject is on the vague border between believing and failing to believe.” Participants attributed knowledge to the protagonists in these cases 77% of the time but attributed belief only 41% of the time. Moreover, in response to an abstract question about the possibility of knowledge without belief, participants were evenly split on whether or not this was possible. Myers-Schulz and Schwitzgebel (forthcoming) conclude:

We thus hope to force those who hold that knowledge entails belief to develop a more substantial argument for their view than the wait-for-counterexamples strategy. Putative counterexamples are available; people’s judgments about them are, we hope to show, divided; and thus a more theoretical approach to the question is necessary, perhaps one that can account for the divided judgments.

When combined with the results of Beebe and Jensen (2012), the results of Experiments 1 through 3 suggest that participants may be more divided on whether belief was present in the

four Knobe effect cases employed above than they were on the question of whether knowledge was present. Experiment 4 tests this hypothesis more directly by asking participants about both knowledge and belief in these cases.

5.1. Methods

5.1.1. Participants

289 undergraduate college students (45% female, predominantly Caucasian, average age = 20) from a large, public university in the northeastern United States participated in this study on a voluntary basis. They were neither paid nor given course credit for their participation.

5.1.2. Materials and procedure

In a mixed experimental design, participants were presented with one of the eight vignettes that have served as the backbone of the experiments above—i.e., either the positive or the negative versions of ENVIRONMENT, MOVIES, SALES, or NAZI. Each participant was then asked to indicate both (a) the extent to which they agreed or disagreed that the protagonist in question believed that his/her primary action would bring about the side-effect in question and (b) the extent to which they agreed or disagreed that the protagonist knew that his/her primary action would bring about the side-effect in question. The order of these two questions was counterbalanced.

5.2. Results

Participants were significantly more inclined to attribute knowledge than belief in ENVIRONMENT and NAZI and were more inclined to attribute belief than knowledge in

SALES (cf. Figure 4 and Table 4).¹⁰ Effect sizes for these factors were medium. Familiar effects for valence were found in ENVIRONMENT, MOVIES, and NAZI and were quite large.¹¹ As in Experiment 1, participants were more inclined to attribute belief in ENVIRONMENT, MOVIES, and NAZI but were evenly split in SALES. Experiment 4 also replicates the findings of Beebe and Jensen (2012), with participants being more inclined to attribute to attribute knowledge in negative cases than in positive ones. The only significant interaction effect between valence and knowledge/belief attribution was found in NAZI.¹²




Figure 4. Mean belief and knowledge attributions in each of the eight conditions of Experiment 4.¹³

¹⁰ Environment: $F(1, 67) = 6.542, p < .05$, partial $\eta^2 = .089$. Movies: $F(1, 72) = 1.150, p > .05$, partial $\eta^2 = .016$. Sales: $F(1, 71) = 6.933, p < .05$, partial $\eta^2 = .089$. Nazi: $F(1, 71) = 5.369, p < .05$, partial $\eta^2 = .070$.

¹¹ Environment: $F(1, 67) = 50.801, p < .001$, partial $\eta^2 = .431$. Movies: $F(1, 72) = 18.220, p < .001$, partial $\eta^2 = .202$. Sales: $F(1, 71) = .822, p > .05$, partial $\eta^2 = .011$. Nazi: $F(1, 71) = 11.088, p < .01$, partial $\eta^2 = .135$.

¹² $F(1, 71) = 4.910, p < .05$, partial $\eta^2 = .065$

¹³ Due to the large number of significant differences, no brackets with asterisks were used to indicate these differences in Figure 4 in order to keep the graph from becoming too cluttered.

	ENVIRONMENT		MOVIES		SALES		NAZI	
Mean	Help	Harm	Better	Worse	Increase	Decrease	Fulfill	Violate
Believed	-.26	1.71	-.17	1.49	.90	.85	.22	.76
Knew	.51	2.26	.54	1.28	-.08	.59	.25	2.00

Table 4. Mean belief and knowledge attributions in each of the eight conditions of Experiment 4.

Across all four cases, 34.6% of participants were more confident that the central protagonists knew the side-effects in question were going to occur than they were that the protagonists believed they would occur (cf. Figure 5). Although this percentage does not differ significantly from the 35.6% who equally attributed knowledge and belief, it is not an insubstantial proportion.¹⁴ In ENVIRONMENT and NAZI, the percentage of participants who were more confident that the central protagonists knew the side-effects in question were going to occur than they were that the protagonists believed they would occur increased to 42.3% (cf. Figure 6). This percentage differs significantly from the 37.3% who equally attributed knowledge and belief in these cases.¹⁵ No gender effects were found.

¹⁴ Two-sample Kolmogorov-Smirnov $D = .0588, p > .05$.

¹⁵ Two-sample Kolmogorov-Smirnov $D = .1690, p < .05$.

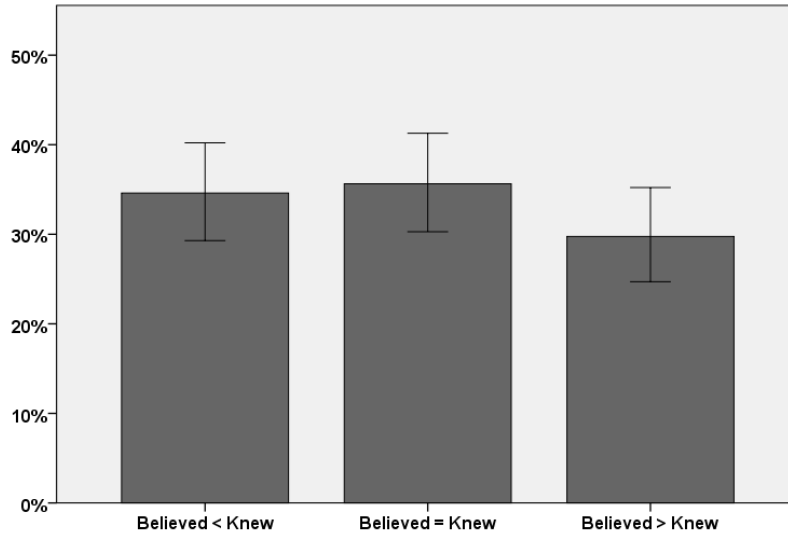


Figure 5. Percentages of participants who were less inclined to attribute belief than knowledge (34.6%), equally inclined to attribute belief and knowledge (35.6%), and more inclined to attribute belief than knowledge (29.8%) across all conditions of Experiment 4.¹⁶




Figure 6. Percentages of participants who were less inclined to attribute belief than knowledge (42.3%), equally inclined to attribute belief and knowledge (37.3%), and more inclined to attribute belief than knowledge (20.4%) in the ENVIRONMENT and NAZI conditions of Experiment 4.¹⁷

¹⁶ I did not represent significant differences from uniformity (i.e., an even 33.3% split across all three categories), since such differences would not be theoretically meaningful.

¹⁷ The percentages in the first two groups also differed significantly from that of the third group, but since

5.3. Discussion

Commenting on their findings, Myers-Schulz and Schwitzgebel (forthcoming) write:

We do not assert that a majority of respondents have intuitions in conformity with the view that knowledge does not entail belief, but only that a substantial proportion do, perhaps about half.

The results of Experiment 4 do not feature differences between participant attributions of knowledge and belief as large as those found by Myers-Schulz and Schwitzgebel (forthcoming). However, Myers-Schulz and Schwitzgebel uncovered their result in cases that were carefully crafted for their anticipated ability to lead ordinary participants to attribute knowledge without belief. The Knobe effect cases employed in the present study were not constructed or selected for this purpose, and yet they reveal a similar effect. There were more participants whose confidence in attributing belief was less than their confidence in attributing knowledge than there were participants whose confidence measures were equal or inclined in the opposite direction. This fact should lend additional credence to Myers-Schulz and Schwitzgebel's suggestion that the view that knowledge entails belief may not be as widely shared among the folk as epistemologists have traditionally thought.

6. Experiment 5

In response to the work of Myers-Schulz and Schwitzgebel (forthcoming), Rose and Schaffer (forthcoming) have argued that knowledge does entail belief but that the relevant sense of belief is “not the *occurrent* notion of a thought consciously endorsed but rather the *dispositional* notion of information available to mind.” Consequently, Rose and Schaffer reran variants of some of Myers-Schulz and Schwitzgebel's vignettes but asked participants whether they agreed that each

these differences are theoretically inconsequential, I did not represent them in Figure 6.

protagonist (a) knew the facts in question and (b1) was disposed to believe those facts or (b2) “still believed” those facts “(in the sense that she still held the information in their mind even if she could not access it).” They found that when a dispositional notion of belief was explicitly triggered, participants attributed both knowledge and belief, thereby seeming to confirm Rose and Schaffer’s hypothesis that knowledge entails dispositional but not occurrent belief. In order to investigate the respects in which elements of the Knobe effect might influence attributions of dispositional belief and to see whether knowledge might entail dispositional belief in these cases, the following three-part experiment was performed.

6.1. Methods

6.1.1. Participants

180 undergraduate college students (67% female, predominantly Caucasian, average age = 22) from a large, public university in the northeastern United States participated in Experiments 5a and 5b by filling out an online questionnaire in exchange for extra credit in one of their college courses. 235 undergraduate college students (48% female, predominantly Caucasian, average age = 21) from a large, public university in the northeastern United States participated in Experiment 5c on a voluntary basis. They were neither paid nor given course credit for their participation.

6.1.2. Materials and procedure

Experiments 5a and 5b each employed a mixed experimental design, in which each participant read either all four positive versions of the vignettes employed in the experiments above or all four of the negative versions. Participants in Experiment 5a were asked the extent to which they agreed or disagreed with the relevant instances of the following schema:

- [The protagonist] was disposed to believe that [the protagonist's primary action] would [lead to the side-effect in question].

Participants in Experiment 5b were asked to indicate the extent to which they agreed or disagreed with the relevant instances of the following schema:

- [The protagonist] believed (in the sense that she/he still held the information in her/his mind even though she/he didn't care about it) that [the protagonist's primary action] would [lead to the side-effect in question].

Participants were given the following answers to choose from: Completely Disagree, Mostly Disagree, Slightly Disagree, Neither Agree nor Disagree, Slightly Agree, Mostly Agree, and Completely Agree.

Experiment 5c employed a between subjects design in which each participant read either the positive or the negative version of one of the vignettes and was asked (i) whether the protagonist knew the protagonist's primary action would lead to the side-effect in question and (ii) whether the protagonist was disposed to believe this would happen. The order of these questions was counterbalanced. Participants indicated their responses on a seven-point scale ranging from -3 to 3.

6.2. Results

Although numeric values did not appear next to the online answer choices in Experiments 5a and 5b, for the sake of uniformity of presentation participant responses have been mapped onto a seven-point scale ranging from -3 to 3. In accord with much of the data from Experiments 1 through 4 and other findings associated with the Knobe effect, participants in Experiment 5a

were significantly more inclined to attribute dispositional belief in negative cases than in positive ones (cf. Figure 7 and Table 5).¹⁸




Figure 7. Mean participant responses to the eight vignettes used in Experiment 5a.

	ENVIRONMENT		MOVIES		SALES		NAZI	
	Help	Harm	Better	Worse	Increase	Decrease	Fulfill	Violate
Mean	.41	1.31	.39	1.18	.32	1.40	.34	1.36
SD	1.72	1.94	1.72	2.15	1.71	1.80	1.78	1.92

Table 5. Mean participant responses and standard deviations for each of the four pairs of conditions in Experiment 5a.

In Experiment 5b, however, there was no main effect for valence, and mean participant responses differed from the neutral midpoint less than in Experiment 5a.¹⁹

¹⁸ $F(1, 84) = 7.488, p < .01, \text{partial } \eta^2 = .082.$




Figure 8. Mean participant responses to the eight vignettes used in Experiment 5b.

	ENVIRONMENT		MOVIES		SALES		NAZI	
	Help	Harm	Better	Worse	Increase	Decrease	Fulfill	Violate
Mean	.50	.70	.37	.98	.74	.88	.22	1.02
SD	1.75	2.22	1.83	2.08	1.78	2.10	1.79	1.98

Table 6. Mean participant responses and standard deviations for each of the four pairs of conditions in Experiment 5b.

In Experiment 5c, participants in six out of the eight conditions were on average more confident that the protagonist knew the side-effect in question would occur than they were that the protagonist was disposed to believe it would occur. However, due to the relatively small sample size (approx. 30 participants per condition), only in NAZI did this difference reach statistical

¹⁹ $F(1, 87) = 1.526, p > .05, \text{partial } \eta^2 = .017.$

significance.²⁰ Because the effect sizes in these conditions were small to medium, it is likely that a significant effect would have been found with a larger sample size. The typically large asymmetry between positively and negatively valenced cases was again found in every condition except SALES.²¹

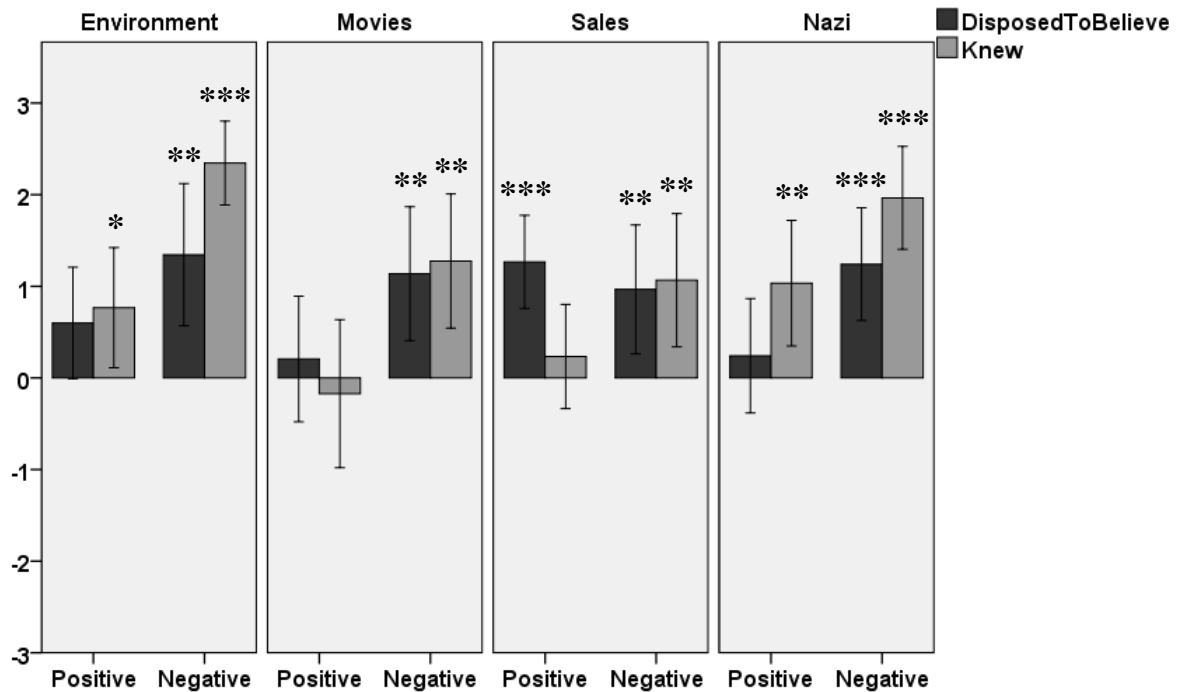


Figure 9. Mean participant responses to the eight vignettes used in Experiment 5c.²²

²⁰ Environment: $F(1, 57) = 3.665, p = .061, \text{partial } \eta^2 = .060$. Movies: $F(1, 56) = .113, p > .05, \text{partial } \eta^2 = .002$. Sales: $F(1, 58) = 2.258, p > .05, \text{partial } \eta^2 = .037$. Nazi $F(1, 56) = 6.178, p < .05, \text{partial } \eta^2 = .099$.

²¹ Environment: $F(1, 57) = 13.532, p < .01, \text{partial } \eta^2 = .192$. Movies: $F(1, 56) = 10.690, p < .01, \text{partial } \eta^2 = .160$. Sales: $F(1, 58) = .744, p > .05, \text{partial } \eta^2 = .013$. Nazi $F(1, 56) = 27.034, p < .01, \text{partial } \eta^2 = .153$.

²² Due to the large number of significant differences, no brackets with asterisks were used to indicate these differences in Figure 9 in order to keep the graph from becoming too cluttered.

	ENVIRONMENT		MOVIES		SALES		NAZI	
Mean	Help	Harm	Better	Worse	Increase	Decrease	Fulfill	Violate
Disposed to Believed	.60	1.34	.21	1.12	1.27	.97	.24	1.24
Knew	.77	2.34	-.17	1.28	.23	1.07	1.03	1.97

Table 7. Mean participant responses for each of the four pairs of conditions in Experiment 5c.

Across all four conditions there were slightly more participants who displayed less confidence about the attribution of dispositional belief than about knowledge than there were participants who were equally confident about both (cf. Figure 10).²³ The proportion of participants who were less confident about dispositional belief than about knowledge was even greater when looking only at ENVIRONMENT and NAZI.²⁴




Figure 10. Percentages of participants who were less inclined to attribute dispositional belief than knowledge (36.6%), equally inclined to attribute dispositional belief and knowledge (31.1%), and more inclined to attribute dispositional belief than knowledge (32.3%) across all conditions of Experiment 5c.

²³ This difference, however, failed to reach statistical significance. Two-sample Kolmogorov-Smirnov $D = .0818$, $p > .05$.

²⁴ Two-sample Kolmogorov-Smirnov $D = .0684$, $p > .05$.




Figure 11. Percentages of participants who were less inclined to attribute dispositional belief than knowledge (42.7%), equally inclined to attribute dispositional belief and knowledge (35.9%), and more inclined to attribute dispositional belief than knowledge (21.4%) in the ENVIRONMENT and NAZI conditions of Experiment 5c.

6.3. Discussion

Despite the fact that Rose and Schaffer (forthcoming) found that participants in their experiments were more likely to attribute knowledge and dispositional belief equally than they were to favor the former over the latter, Experiment 5 revealed a different pattern. One difference between the experimental designs of Rose and Schaffer’s experiments and the present one is that Rose and Schaffer followed Myers-Schulz and Schwitzgebel (forthcoming) in using a forced choice answer format that prompted participants to choose between ‘Yes’ and ‘No,’ whereas Experiment 5 employed a Likert scale. Because the vignettes employed in Experiments 4 and 5—unlike those of Myers-Schulz and Schwitzgebel—were not specially crafted for the purpose of undermining the view that knowledge entails belief, the results of Experiments 4 and 5 seem

to provide enough support to Myers-Schulz and Schwitzgebel's challenge to keep it alive in the face of Rose and Schaffer's attempt to put it to rest.

7. General discussion

Pettit and Knobe (2009) suggest that our practices of normative evaluation appear to have an utterly pervasive effect on folk psychological attributions. Data from the experiments described above provide further confirmation for this view. Participants were more likely to ascribe belief, higher degrees of belief, higher degrees of rational belief, and dispositional belief to agents in central Knobe effect cases who bring about negative side-effects than to agents who bring about positive ones.

Knobe (2010, p. 317) has claimed that the contrasting help and harm conditions in ENVIRONMENT (and by implication other Knobe effect cases as well) "are exactly the same in almost every respect but differ in their moral status." The present set of experiments, together with the work of Beebe and Buckwalter (2010) and Beebe and Jensen (2012), however, shows that naïve participants take the protagonists in these conditions to differ psychologically in important ways.

Most scholars working on the Knobe effect share Knobe's assumptions about how the mental states of the chairman and other protagonists will be viewed, and this in turn shapes the explanations they offer of the Knobe effect data. Mele and Cushman (2007), for example, dub the target actions of protagonists in Knobe effect cases as 'ASEAs' or 'anticipated side-effect actions,' and Cushman and Mele (2007) refer to all chairman-style cases as 'B+D-' cases—i.e., as cases where the protagonist believes but does not desire that a certain side-effect will occur. Commenting upon Nichols and Ulatowski's (2007) finding that approximately one third of

participants in a within-subjects experiment on the Knobe effect attribute intentionality in both the help and the harm conditions of ENVIRONMENT and another third attribute intentionality in neither, Cushman and Mele (2007, p. 173) write:

The fact that certain people judge both CEOs to have intentionally performed the actions asked about... whereas other people judge both CEOs to have unintentionally performed those actions... is evidence that people hold differing views about the sufficiency of belief for intentional action and the necessity of desire.

Cushman and Mele (2007, p. 175) conjecture that participants who attribute intentionality in both conditions employ a concept of intentional action that is characterized by the following rule: “An action is intentional if it is performed with belief.” They also hypothesize that participants who give asymmetrical responses operate with a concept that obeys the following rule: “Belief is a sufficient condition for intentional action in and only in cases of *morally bad* actions.” However, the data that Cushman and Mele cite fail to serve as evidence for any of these claims because participants do not in fact attribute belief equally in both conditions. This means that the case for Cushman and Mele’s proposal that there are ‘two and a half’ folk concepts of intentional action rests upon a mistake. Similar considerations apply to at least part of Lanteri’s (2012) ‘three and a half concepts’ approach, which builds upon the work of Cushman and Mele (2007).

The last few years have witnessed an explosion in the number of hypotheses being put forward to explain the original Knobe effect. Too few of these explanations, however, are empirically motivated or attempt to uncover new empirical facts about how participants respond to and think about the central cases in the literature. By experimentally examining four Knobe effect cases from a variety of angles, it is hoped that the theoretical debate concerning the proper explanation of the effect will be enriched and illuminated.

References

- Beebe, J. R., and W. Buckwalter. 2010. The epistemic side-effect effect. *Mind & Language* 25: 474-498.
- Beebe, J. R., and M. Jensen. 2012. Surprising connections between knowledge and action: The robustness of the epistemic side-effect effect. *Philosophical Psychology* 25: 689-715.
- Cova, F., and H. Naar. forthcoming. Side-effect effect without side effect: The pervasive impact of moral considerations on judgments of intentionality. *Philosophical Psychology*.
- Cushman, F., and A. Mele. 2007. Intentional action: Two-and-a-half folk concepts? In J. Knobe and S. Nichols (Eds.), *Experimental Philosophy* (pp. 171-188). New York: Oxford University Press.
- Doris, J., J. Knobe, and R. L. Woolfolk. 2007. Variantism about responsibility. *Philosophical Perspectives* 21: 183-214.
- Hindriks, F. 2008. Intentional action and the praise-blame asymmetry. *Philosophical Quarterly* 58: 630-641.
- Hsee, C. K. 1998. Less is better: When low-value options are valued more highly than high-value options. *Journal of Behavioral Decision Making* 11: 107-21.
- Knobe, J. 2003a. Intentional action and side effects in ordinary language. *Analysis* 63: 190-3.
- Knobe, J. 2003b. Intentional action in folk psychology: An experimental investigation. *Philosophical Psychology* 16: 309-24.
- Knobe, J. 2004a. Folk psychology and folk morality: Response to critics. *Journal of Theoretical and Philosophical Psychology* 24: 270-279.
- Knobe, J. 2004b. Intention, intentional action and moral considerations. *Analysis* 64: 181-187.

- Knobe, J. 2006. The concept of intentional action: A case study in the uses of folk psychology. *Philosophical Studies* 130: 203-231.
- Knobe, J. 2007. Reason explanation in folk psychology. *Midwest Studies in Philosophy* 31: 90-106.
- Knobe, J. 2010. Person as scientist, person as moralist. *Behavioral and Brain Sciences* 33: 315-329.
- Knobe, J., and A. Burra. 2006. The folk concepts of intention and intentional action: A cross-cultural study. *Journal of Cognition and Culture* 6: 113-132.
- Knobe, J., and J. Doris. 2010. Responsibility. In J. Doris and the Moral Psychology Research Group (Eds.), *The Handbook of Moral Psychology*. Oxford: Oxford University Press, pp. 321-54..
- Knobe, J., and G. Mendlow. 2004. The good, the bad and the blameworthy. *Journal of Theoretical and Philosophical Psychology* 24: 252-258.
- Lanteri, A. 2012. Three-and-a-half folk concepts of intentional action. *Philosophical Studies* 158: 17-30.
- Leslie, A. M., J. Knobe, and A. Cohen. 2006. Acting intentionally and the side-effect effect: Theory of mind and moral judgment. *Psychological Science* 17: 421-427.
- Machery, E. 2008. The folk concept of intentional action: Philosophical and experimental issues. *Mind & Language* 23: 165-189.
- McCann, H. 2005. Intentional action and intending: Recent empirical studies. *Philosophical Psychology* 18: 737-748.
- Mele, A. 2006. The folk concept of intentional action: A commentary. *Journal of Cognition and Culture* 6: 277-290.

- Mele, A., and F. Cushman. 2007. Intentional action, folk judgments, and stories: Sorting things out. *Midwest Studies in Philosophy* 31: 184-201.
- Myers-Schulz, B., and E. Schwitzgebel. forthcoming. Knowing that p without believing that p. *Noûs*.
- Nadelhoffer, T. 2004a. On praise, side effects, and folk ascriptions of intentionality. *Journal of Theoretical and Philosophical Psychology* 24: 196-213.
- Nadelhoffer, T. 2004b. Blame, badness, and intentional action: A reply to Knobe and Mendlow. *Journal of Theoretical and Philosophical Psychology* 24: 259-269.
- Nadelhoffer, T. 2006. Bad acts, blameworthy agents, and intentional actions: Some problems for jury impartiality. *Philosophical Explorations* 9: 203-220.
- Nichols, S., and J. Ulatowski. 2007. Intuitions and individual differences: The Knobe effect revisited. *Mind & Language* 22: 346-365.
- Pettit, D., and J. Knobe. 2009. The pervasive impact of moral judgment. *Mind & Language* 24: 586-604.
- Phelan, M. forthcoming. Evidence that stakes don't matter to evidence. *Philosophical Psychology*.
- Phelan, M., and H. Sarkissian. 2008. The folk strike back: Or, why you didn't do it intentionally, though it was bad and you knew it. *Philosophical Studies* 138: 291-298.
- Rose, D., and J. Schaffer. forthcoming. Knowledge entails dispositional belief. *Philosophical Studies*.
- Sverdlik, S. 2004. Intentionality and moral judgments in commonsense thought about action. *Journal of Theoretical and Philosophical Psychology* 24: 224-236.

Tannenbaum, D., P. H. Ditto, and D. A. Pizarro. 2007. Different moral values produce different judgments of intentional action. Unpublished manuscript, University of California-Irvine.