JOURNAL OF COGNITION AND CULTURE 16 (2016) 1 - 28

BRILL

brill.com/jocc

Evidence for the Context Dependence of the Side-Effect Effect

*Stephan Lau**
Rainer Reisenzein
Institute for Psychology, Ernst Moritz Arndt University of Greifswald, Friedrich-Ludwig-Jahn-Straße 17a, 17487 Greifswald, Germany
**Corresponding author, e-mail: stephan.lau@uni-greifswald.de*

Abstract

In four experiments involving 565 German speakers we tested hypotheses about possible determinants of the side-effect effect (SEE), which consists of judging foreseen bad, but not good, side-effects of actions as intentionally produced. Experiment 1 failed to find intentionality ascriptions for bad side-effects for the majority of the participants in two different scenarios and obtained no consistent support for two hypothesized social-cognitive determinants of the SEE, the agent's attitude and the mode of effect description. Experiment 2 replicated the SEE in the original CEO scenario, but again found no evidence that the effect was influenced by the agent's attitude towards the side-effect. The SEE was also not influenced by a manipulation of the moral quality of the agent's primary goal. Experiment 3 investigated six additional scenarios used in previous studies and again obtained clear evidence for the SEE only in the CEO scenario. In addition, Experiment 3 demonstrated that judgments of both intent and intentionality strongly increased if the original side-effect was described as a means to the agent's primary goal, or as an independently pursued goal. Taken together, the findings suggest that for German speakers, the SEE depends on the specifics of the scenario content and is difficult to obtain outside the original CEO scenario. Consistent with these conclusions, Experiment 4 documented parallel difficulties replicating the "means effect", an analogue of the SEE on the level of means, but replicated the SEE in a scenario closely modeled after the original CEO scenario.

Keywords

 | DOI 10.1163/15685373-12342180

intentionality judgments – moral reasoning – scenario dependence of the side-effect effect – side-effect effect

1 Introduction

Traditional theories of moral judgment assume that the attribution of intentionality is a precursor to moral evaluation (e.g., Meinong, 1894; Heider, 1958; Fincham and Jaspars, 1980; Weiner, 1985; McGraw, 1987). For example, according to an influential model of responsibility attribution proposed by Heider (1958), intentionality is a central cue to responsibility, which in turn is a presupposition of moral blame and praise. Similar views are held by more recent authors (e.g., Weiner, 1995; Guglielmo and Malle, 2010). Although these theories of moral judgment make intuitive sense and can claim empirical support (e.g., Fincham and Jaspars, 1980; Weiner, 1995; Malle and Knobe, 1997; Skitka et al., 2005), recent work on intentionality ascriptions for side-effects has cast doubt on their generality (Knobe 2003, 2006). Knobe (2003) presented participants with a scenario where a CEO's decision to start a new program had either a foreseen bad or good side-effect (it harmed or helped the environment), and asked them whether or not the side-effect had been brought about intentionally by the agent. Although only 23% of the participants said that the agent intentionally produced the good side-effect, 82% said that he intentionally produced the bad side-effect. Called the side-effect effect (SEE; Leslie et al., 2006), this finding has since been replicated in a variety of scenarios in addition to the original CEO scenario (e.g., Knobe and Mendlow, 2004; Knobe and Burra, 2006; Malle, 2006; Nadelhoffer, 2006a; Phelan and Sarkissian, 2008; Pellizzoni et al., 2009; see Experiment 3 for additional references).

The SEE has attracted interest for at least two reasons. First, it challenges the traditional model of responsibility judgment and moral evaluation (e.g., Heider, 1958), as well as a traditional analysis of intentionality judgments (the so-called "simple view"), according to which an action is regarded as intentional only if it is intended (Bratman, 1984). The SEE seems to reveal a moral bias in intentionality judgments that is incompatible with these "rational" theories of the attribution of intentionality and responsibility, and threatens to undermine the fairness of the moral evaluations and actions (such as punishment) that may result from them (Nadelhoffer, 2006b). Second, the SEE seems to demonstrate that analyses of folk-psychological concepts proposed by philosophers and psychologists can significantly deviate from the actual meanings of these concepts in ordinary language, supporting the argument that "arm-chair" conceptual analyses should be checked by empirical studies of

concept use (e.g., Reisenzein, 1995; Nadelhoffer and Nahmias, 2007).

Because of the theoretical and practical interest that attaches to the SEE, considerable effort has been exerted to explain it (for recent reviews, see Cova et al., 2012; Sloman et al., 2012). The causes of the SEE invoked in the proposed explanations range form a moral asymmetry in intentionality judgments (e.g., Knobe, 2006) to the informative value of norm-violations (e.g., Uttich and Lombrozo, 2010; see also Holton, 2010), to linguistic (e.g., Cova et al., 2012) and methodological factors (e.g., Guglielmo and Malle, 2010). However, so far, no generally accepted explanation of the SEE has been found.

The original goal of the present research was to replicate the SEE for German speakers and to test a number of hypotheses about possible determinants of this effect (Experiment 1). Unexpectedly, however, the majority of the participants did not attribute intentionality for negative side-effects, and hence there was no clear evidence for the SEE. As mentioned, the SEE is defined the as tendency to judge foreseen, morally bad side-effects of actions, but not good side-effects, as intentionally produced (e.g., Nichols and Ulatowski, 2007; Machery, 2008). On the group level, on which the SEE is usually studied, this tendency manifests itself in a high percentage of intentionality attributions for bad side-effects, coupled with a low percentage of intentionality attributions for good side-effects (e.g., Knobe, 2003; Nichols and Ulatowski, 2007). In Experiments 1–3 we took the previous results concerning good side-effects for granted and therefore included only bad side-effect scenarios. For these scenarios, we took an SEE to be present if at least the majority (>50%) of the participants judged the side-effect as intentionally produced. Although still weaker definitions of the SEE are possible (in particular: the percent of intentionality ascriptions for bad side-effects is greater than that for positive side-effects), they are considerably less interesting, because they imply that the SEE is a minority phenomenon.Because the original CEO scenario (Knobe, 2003) had not been included in Experiment 1, a second study was conducted that focused on this scenario, while again testing hypotheses about possible determinants of the SEE. This time, attributions of intentionality for the bad side-effect were obtained for the majority of the participants, although no support was again found for the tested determinants of the SEE. This led us to hypothesize that the SEE is, at least for German speakers, a context-dependent effect, that is, its occurrence depends on the specifics of the social situation described in the scenario — the specific actor, his or her social role, the specific action at issue, the specific effects of the action etc. To explore this hypothesis further, we studied intentionality ascriptions for unintended side-effects in six additional scenarios taken from the SEE literature (Experiment 3) and in a new

scenario that closely resembled the CEO scenario without duplicating it (Experiment 4). In addition, Experiment 4 examined whether the difficulties of obtaining an SEE for German speakers generalized to the so-called "means effect" recently reported by Cova and Naar (2012), an analogue to the SEE on the level of means.

2 Experiment 1

The aim of Experiment 1 was to replicate the SEE for native speakers of German and to test two hypotheses about determinants of the SEE that seemed to have some prima facie plausibility: the agent's attitude towards the side-effect, and the mode of description of the action effects. For reasons explained in Footnote 1, only negative side-effect scenarios were included in this study (as well as in Experiments 2 and 3).

2.1 *Agent's Attitude towards the Side-Effect*

We reasoned (with Guglielmo and Malle, 2010) that the attitude of the agent towards the side-effect should be an important piece of information for judgments of intentionality. If the agent has a positive attitude towards the side-effect, one can infer that the agent desired this effect to occur and performed the action, at least in part, to achieve it. It would then not be inappropriate to conclude that the agent produced the side-effect intentionally. In contrast, if the agent's attitude towards the side-effect is negative, one can infer that the agent did not want this effect to occur and therefore did not intend to produce it. In this case, it would be appropriate to conclude that the agent did not produce the side-effect intentionally.

An inferred positive attitude of the agent towards the side-effect could have contributed to the high percentage of intentionality judgments reported by Knobe (2003). Although the agent in Knobe's CEO scenario did not explicitly approve of the side-effect, he expressed a non-caring attitude towards it ("I don't care at all about harming the environment"; Knobe, 2003:191) that, it can be argued, revealed at least a minimal degree of approval. Supporting this assumption, Phelan and Sarkissian (2008) found (although with a different scenario) that if an agent said that he felt "terrible about increasing joblessness", the SEE effect was strongly reduced (attributions of intentionality dropped below 50%).

To test the hypothesis that the SEE depends on the agent's attitude towards the side-effect, we varied the verbal expression of the attitude similar to Guglielmo and Malle (2010): The agent either expressed his approval of the side-

effect (positive attitude condition), his regret about the side-effect (negative attitude condition), or did not reveal his attitude (control condition). Additionally, we included a condition in which the agent expressed a "don't care" attitude, as in Knobe's (2003) original study.

2.2 *Mode of Description of Action Effects*

The second potential determinant of the SEE studied in Experiment 1 concerned the description of the action effects: In one condition, the participants were only informed that the bad side-effect had occurred, whereas in a second condition, they were explicitly informed that both the side-effect and the intended effect had occurred. Previous studies did not systematically control for this factor, but it could conceivably have played a role for the obtained results. For example, in Knobe's (2003) CEO scenario, only the occurrence of the bad side-effect was mentioned and an SEE was found for most participants, whereas Phelan and Sarkissian (2008) informed their participants (although again using a different scenario) that both the intended and the unintended effect occurred, and found that the percentage of intentionality judgments was reduced to 29%.

2.3 *Method*

2.3.1 Participants

The sample consisted of 320 students (99 males and 219 females, 2 did not indicate their sex) at the University of Greifswald, between 18 and 29 years of age. Potential participants were contacted through a student mailing list and were provided with the link to the webpage of the experiment. Participation was voluntary, with no extra reward. About 60 additional participants were excluded from the data analyses because they indicated on the start page of the experiment that they just wanted to "take a look" rather than to participate seriously, because they did not complete all pages of the experiment, or because they processed pages repeatedly (using the "Back" button of the web browser).

2.3.2 Scenarios

Two different scenarios were used. The first, which was an adaptation of Knobe's (2003) scenario to student life, described a minister of education who decides to implement an educational policy. The second scenario (adapted from Nadelhoffer, 2006a; case C6) described a military sniper who decides to shoot at his target. This scenario was included because intentionality ascriptions for the negative side-effect were comparatively low (55%), leaving more

room for a possible effect of positive agent attitude to emerge.

In the first scenario, a minister of education decides to implement an educational policy that will have the negative side-effect of reducing the number of available study places. In its basic form (see below for experimental variations), the scenario read as follows:

> The Minister of Education of a state is visited by his State Secretary with an important issue. "We have developed a new educational program for our universities. It will markedly increase the quality of study and research. However, many applicants will no longer get a place to study."
> The minister thinks: "Hm…it is my goal to bring study and research to a high standard of excellence. Let's start the program."
> The program was started. Sure enough, many applicants got no place to study.

In the second scenario, a sniper decides to shoot at his target, knowing that by shooting he will uncover his position to the enemy (Nadelhoffer, 2006a). In its basic form, this scenario (which was transferred by us into a World War II context) read as follows:

> During World War II, a sniper has the mission to kill a high officer of the enemy to cover his comrades' retreat. He has detected his target in the backcountry and now faces a problem. If he shoots, he will hit his target, but will also uncover his position through the muzzle flash and thus get spotted.
> The sniper thinks. "It is my aim to complete the mission. I will shoot."
> The sniper shoots. Sure enough, he was spotted through the muzzle flash.

2.3.3 Design

The experiment was based on a 4 (attitude of the agent towards the side-effect) × 2 (mode of description of the action effects) between-subjects design. The attitude of the agent towards the side-effect was manipulated by including different verbal expressions of his attitude into the scenario description (here illustrated for the minister scenario): A–, negative attitude: in this condition, the agent openly regrets the side-effect ("… although I find it terrible if fewer applicants get a place to study"); A0, no attitude expressed; A~, don't care attitude: this condition used Knobe's (2003) wording ("… and I don't care if fewer applicants will get a place to study"); A+, positive attitude: in this condition, the agent publicly approves of the side-effect ("… I prefer to have fewer students who can study better"). Analogous wordings were used for the attitude

expressions of the agent in the sniper scenario. For example, in the positive attitude condition, the sniper muses "It is my aim to complete the mission and to achieve this goal, being uncovered is a price I am very happy to pay".

The description of the effects of the agent's action was varied as follows: In the side-effect only condition (SE), only the occurrence of the negative side-effect was mentioned (e.g., "Sure enough, many applicants got no place to study"). This corresponds to the effect description in the original CEO scenario (Knobe, 2003). In contrast, in the side-effect plus intended effect condition (SE+IE), both the intended effect and the unintended negative side-effect were explicitly said to have occurred (e.g., "Sure enough, the quality of study and research increased, but many applicants got no place to study").

To avoid systematic transfer effects, each of the eight resulting conditions of the minister scenario was paired with a maximally different (in terms of the experimental manipulations) version of the sniper scenario (e.g., condition A+/SE was paired with condition A-/SE+IE) and the 320 participants were randomly assigned to these scenario combinations. Hence, each participant first judged one of the eight versions of the minister scenario, and then a very dissimilar version of the sniper scenario. After all questions accompanying both scenarios had been answered, the participants were debriefed about the purpose of the study.

2.3.4 Dependent Variables

The two central dependent variables were judgments of perceived intentionality and intent regarding the negative side-effect. Intentionality was assessed with two different items. The first, dichotomous item corresponds to the intentionality question used in most previous studies of the SEE. For the minister scenario, this item was: "Did the minister intentionally (German: 'absichtlich') see to it that many applicants got no study place?" (yes/no). The second intentionality item asked the same question, but was to be answered on an 11-point rating scale with endpoints "No" and "Yes". The dichotomous item was presented immediately after the description of the scenario, whereas the rating was presented as the last question. The rating was included to validate the results obtained for the dichotomous item and to capture variance in beliefs about intentionality that might be missed by a forced-choice response.

The intention of the agent to produce the bad side-effect was assessed with the question: "Was it the minister's intention (German: 'Absicht') that many applicants got no study place?" (yes/no). Previous SEE research has found that there is a gap between endorsements of intentionality (present) and intention (absent) regarding bad side-effects (e.g., McCann, 2005; Knobe and Burra, 2006; Nadelhoffer, 2006a), a finding that seems to support the claim that the

"simple view" of intentionality judgments (Bratman, 1984) is invalid in moral contexts.

Two additional items were included to verify that the scenarios were perceived as intended. The first item (badness of the side-effect) asked for the participant's own moral evaluation of the side-effect (e.g. in the minister scenario: "How do you judge the fact that many applicants got no place to study?"), whereas the second item (agent approval) asked for the perceived attitude of the agent towards the side-effect (e.g., "How bad do you think does the minister regard the fact that many applicants got no place to study?"). Both items were answered using 11-point rating scales with endpoints labeled not bad at all and very bad.

2.3.5 Procedure

The study was implemented as a web experiment using WEXTOR (http://wextor.org/wextor/en/; Reips and Neuhaus, 2002) and was run on the experiment web server of the Institute of Psychology. On the start page of the experiment, the participants were informed that the study was concerned with "how people make judgments about the actions of others" and were asked to indicate whether they wanted to take part in the study seriously, or to "just take a look". On the next page, after indicating their age and sex, the participants were told that they would be asked to judge the actions of two agents in two different scenarios, and that there were no true or false answers, but their personal opinion mattered. Subsequently, the first scenario was presented, followed by the five questions. Each item was presented on a separate HTML page. Item order was fixed, with the dichotomous intentionality and intention items presented first, followed by a refresh of the scenario, the manipulation checks, and the intentionality rating.

2.4 *Results*

2.4.1 Manipulation Checks

As intended, the side-effect was regarded as bad in all experimental conditions of both the minister scenario (M = 7.02; SD = 2.21) and the sniper scenario (M = 6.58, SD = 2.89) (see Table 1). Analyses of variance (ANOVAs) revealed no significant effects of the experimental manipulations on the perceived badness of the side-effect in either scenario, p values ≥ 0.18.

Also as intended, the rating of the agent's approval of the side-effect revealed a significant main effect of his expressed attitude towards the side-effect in both the minister scenario, $F(3, 312) = 24.0$, $p < 0.001$, $\eta^2 = 0.18$, and the sniper scenario, $F(3, 312) = 6.1$, $p < 0.001$, $\eta^2 = 0.06$. The mode of description of

the action effects had no significant effect on the agent approval rating, and the interaction of the two factors was likewise nonsignificant (p values ≥0.314). As expected, in both scenarios, participants believed that the agent increasingly approved of the side-effect, the more positive his expressed attitude was (Table 1). Hence, the attitude manipulation was successful.

TABLE 1 *Descriptive Statistics, Experiment 1*

	Condition							
	SE				SE + IE			
	A–	A0	A~	A+	A–	A0	A~	A+
Minister scenario								
Dichotomous items (percentage of Yes answers)								
Intentionality	17.5	20.0	37.5	30.0	32.5	30.0	32.5	40.0
Intention	0.0	2.5	12.5	22.5	10.0	7.5	10.0	15.0
Rating items (*M*, SD)								
Intentionality	3.10 (2.89)	3.35 (2.95)	3.82 (3.53)	4.50 (3.19)	3.84 (3.33)	3.10 (2.72)	3.85 (3.06)	4.48 (2.83)
Badness of side-effect	7.30 (1.96)	7.05 (2.26)	7.45 (1.71)	6.42 (2.53)	6.52 (2.35)	7.55 (1.91)	7.02 (2.46)	6.82 (2.26)
Approval	3.92 (2.12)	3.38 (1.86)	1.55 (1.50)	2.30 (1.54)	4.38 (2.10)	3.25 (2.07)	2.28 (1.80)	2.10 (1.99)
Sniper scenario								
Dichotomous items (percentage of Yes answers)								
Intentionality	32.5	30.0	27.5	42.5	30.0	40.0	30.0	35.0
Intention	2.5	0.0	7.5	5.0	0	5.0	5.0	5.0
Rating items (*M*, SD)								
Intentionality	2.90 (2.82)	2.95 (3.62)	2.13 (2.82)	4.00 (3.71)	3.25 (3.26)	3.70 (3.65)	3.40 (2.98)	2.75 (3.42)
Badness of side-effect	7.27 (2.56)	6.42 (3.34)	5.82 (3.10)	6.80 (3.07)	6.95 (2.53)	5.98 (2.84)	6.30 (2.97)	7.12 (2.58)
Approval	7.82 (2.85)	8.02 (2.40)	7.65 (2.62)	6.32 (3.13)	8.27 (2.23)	8.02 (2.40)	7.60 (2.41)	7.05 (2.90)

Abbreviations: A, agent's attitude toward the bad side effect (–, negative; 0, not mentioned; ~, don't care; +, positive); SE, side effect only; SE + IE, side effect and intended effect. Approval: How bad did the agent regard the side effect?

2.4.2 Intentionality

In the minister of education scenario, the overall percentage of "intentional" answers to the dichotomous intentionality question was 30% (see Table 1). A logistic regression analysis using the factors agent attitude and description of the action effects revealed no significant effects (p values ≥ 0.194). Similar findings were obtained for the intentionality rating (which, as expected, correlated fairly strongly with the dichotomous judgment, r = 0.58, p < 0.001): In all experimental conditions, the average intentionality rating of the side-effect was below the scale midpoint (5), overall M = 3.76 (SD = 3.07), with the highest means obtained in the positive attitude conditions (Table 1). An ANOVA revealed a marginally significant effect of agent attitude, $F(3, 312) = 2.55$, p = 0.056. This suggests that the rating was a more sensitive measure of intentionality than the dichotomous intentionality item. However, the effect size was quite small, η^2 = 0.02.

The results for the sniper scenario were similar: The overall proportion of "intentional" judgments was 33%, the average intentionality ratings were below the scale midpoint in all groups (Table 1), and the experimental manipulations had no significant effects on the dichotomous intentionality judgment (Wald z < 1). In this case, the effects of the experimental manipulations on the intentionality rating were also nonsignificant (F values < 1).

2.4.3 Intention

As can be seen from Table 1, the percentage of participants who said that the minister of education intended to bring about the negative side-effect was very low (10% overall). The logistic regression analysis revealed a significant effect of expressed agent attitude (Wald z = 10.13, p = 0.001, odds ratio = 0.46), and a significant interaction between agent attitude and description of the action effects (Wald z = 5.42, p = 0.02, odds ratio = 1.76), whereas the main effect of the description of the action effect was not significant (Wald z = 2.13, p = 0.144). As Table 1 shows, ascriptions of intention were higher if the agent's attitude towards the bad side-effect was positive or if he did not care, than if his attitude was negative or if he did not reveal his attitude. The significant interaction reflects the finding that this effect was less pronounced if the occurrence of both the side-effect and the intended effect were explicitly mentioned in the scenario description.

In the sniper scenario, the percentage of participants who believed that the sniper intended to produce the side-effect was close to zero (3.75% overall). The logistic regression revealed no significant effects of the experimental manipulations (see Table 1).

2.5 *Discussion*

Experiment 1 investigated intentionality ascriptions for bad side-effects in a large sample of native speakers of German using two different scenarios, one adopted from previous SEE research (the sniper scenario, Nadelhoffer, 2006a) and the other newly created (the minister of education scenario). The experiment yielded two unexpected findings. The first unexpected finding was the failure to obtain clear evidence for an SEE (see above): In both scenarios, the bad side-effect was judged as intentional only by a minority (30 and 33%, respectively). In contrast, in Knobe's (2003) study, 82% judged that the CEO intentionally harmed the environment, and in the sniper scenario investigated by Nadelhoffer (2006a), 55% judged that the sniper intentionally alerted the enemy. Both differences to our scenarios are highly significant statistically: CEO vs. minister, χ^2 (N = 398, df = 1) = 68.5, p < 0.001; sniper vs. sniper, χ^2 (N = 440, df = 1) = 16.6, p < 0.001.

The second unexpected finding of Experiment 1 was that the item typically used to assess perceived intentionality in previous SEE studies (the "yes/no" item) was not significantly influenced by the experimental manipulations of agent attitude and effect description in both scenarios. The presumably more sensitive intentionality rating did reveal a marginally significant effect of agent attitude in the minister of education scenario, but the effect was quite small and was not replicated in the sniper scenario. We conclude from these findings that effects of the experimental manipulations, if they exist, are negligible.

Given the unexpected results of Experiment 1, we decided to conduct a second study that used a German translation of Knobe's (2003) original CEO scenario, in which — different from the novel minister scenario and the sniper scenario used in Experiment 1 — a strong SEE has been repeatedly found. Again, two potential determinants of the SEE were experimentally manipulated in Experiment 2. The first was the agents' attitude towards the side-effect already examined in Experiment 1. The second was a factor that conceivably could have been responsible for the low frequency of intentionality attributions obtained in Experiment1: the moral quality of the agent's primary goal. In our scenarios, the primary goal of the agent had positive moral value (i.e., improving study and research; covering the comrades' retreat), whereas in the original CEO scenario, the primary goal of the agent (increasing profit) was morally ambiguous and possibly even regarded as negative by the participants. The moral goodness of the primary goal of the protagonist of our scenarios could have raised the threshold for judging the bad side-effect as intentional (see also Knobe, 2006).

3 Experiment 2

3.1 *Method*

3.1.1 Participants

The participants were 123 students (92 females) at the University of Greifswald, aged between 18 and 45 years (M = 23.8, SD = 4.2). Participants were recruited from the Institute of Psychology's web panel, a data bank of students interested in participating in psychological experiments. Potential participants were informed about the experiment via e-mail and were provided with the link to its webpage. As a reward, they were given the opportunity to participate in a lottery of five Amazon vouchers worth 20 Euros each. Participants were debriefed about the purpose of the study on the last page of the experiment. About 30 additional participants were excluded from the data analyses for the same reasons as in Experiment 1.

3.1.2 Design

Experiment 2 was based on a 3 (agent attitude) × 2 (moral quality of the agent's primary goal) within-subjects design. One of the six cells, condition A~/I– (i.e., a "don't care" attitude plus a negative primary goal, see below), corresponds to the original version of the CEO scenario (Knobe, 2003). This scenario was translated as closely as possible into German. The remaining five versions of the scenario were obtained by changing the basic version to reflect the different levels of agent attitude and moral quality of the primary goal.

Because the results of Experiment 1 suggested that the "positive" and "don't care" attitude of the agent towards the bad side-effect were largely equivalent, only three levels of agent attitude were used in Experiment 2: Negative attitude (A-), no attitude expressed (A0) and don't care (A~). The moral quality of the agent's goal was manipulated as follows: In the negative moral quality condition, the primary goal of the CEO (making as much profit as he could) was described as in the original version of the CEO scenario (Knobe, 2003). In the positive moral quality condition, the aim of the program endorsed by the CEO was described as "making the company's products cheaper and more durable for consumers". The negative side-effect was the same in all conditions (harming the environment).

3.1.3 Procedure and Dependent Variables

The study was again realized as a web experiment. The procedure was very similar to that of Experiment 1, the main difference being that only one scenario was judged by each participant. The dependent variables were the same

as in Experiment 1 plus the following check for the manipulation of the moral quality of the agent's primary goal: "How do you judge the intention of the CEO to make as much profit as he can/to make products cheaper and more durable?". This item was answered on a 21-point bipolar scale ranging from -10 = absolutely blameworthy, through 0 = neither/nor, to +10 = absolutely praiseworthy. The manipulation check item was presented last; the order of the remaining items was the same as in Experiment 1.

3.2 *Results*

3.2.1 Manipulation Checks

As expected, harming the environment was regarded as very bad (overall M = 8.5, SD = 1.7), without significant differences between experimental conditions (p values for both main effects and the interaction ≥0.354). The rating of agent approval revealed a significant effect of agent attitude, $F(2, 117) = 30.8$, $p < 0.001$, $\eta^2 = 0.34$, and the rating of the perceived moral quality of the agent's primary goal revealed a significant effect of the moral quality manipulation, $F(1, 117) = 36.3$, $p < 0.001$, $\eta^2 = 0.23$. As in Experiment 1, participants said that the CEO disapproved of the bad side-effect more strongly, if his expressed attitude towards the side-effect was negative (compared to the "don't care" and the no expressed attitude conditions). Furthermore, as intended, the participants judged the agent's intention to make products cheaper as more praiseworthy than the intention to make profit, which was regarded as slightly negative (Table 2). Both effects were not qualified by interactions, p values ≥ 0.215.

3.2.2 Intentionality

Although the experimental manipulations were successful, their effects on intentionality did not reach conventional levels of significance, neither for the rating (p values ≥ 0.311), nor for the dichotomous item (p values ≥ 0.075). Furthermore, counter to the hypothesis, proposed as a possible explanation of the results of Experiment 1, that a morally good primary goal reduces attributions of intentionality for negative side-effects, these attributions tended to be (albeit nonsignificantly) more frequent if the agent's primary goal was morally good (Table 2).

However, the most important finding of Experiment 2 was that, in agreement with Knobe's (2003) original finding, the production of the side-effect in the CEO scenario was judged as intentional by the clear majority of the participants (overall 76.4%, as compared to 83% in Knobe's, 2003 study). This finding was replicated for the rating scale measure of intentionality (overall M = 6.5, SD = 2.9; correlation with the dichotomous item r = 0.72).

TABLE 2 *Descriptive Statistics, Experiment 2*

Attitude	Positive primary goal			Negative primary goal		
	A–	A0	A~	A–	A0	A~
Dichotomous items (percentage of Yes answers)						
Intentionality	85.0	90.0	75.0	81.8	62.0	65.0
Intention	15.0	65.0	25.0	22.7	28.6	35.0
Rating items (*M*, SD)						
Intentionality	7.0 (2.22)	7.1 (3.14)	6.25 (3.32)	6.68 (2.34)	6.43 (3.25)	5.65 (2.93)
Badness of side-effect	8.30 (1.72)	8.60 (1.98)	8.85 (1.04)	8.09 (1.85)	8.57 (1.86)	8.60 (1.70)
Approval	4.40 (2.30)	2.50 (2.28)	0.65 (0.81)	3.95 (2.01)	1.95 (1.63)	1.40 (1.54)
Moral quality	4.55 (3.36)	5.05 (4.84)	2.55 (5.84)	0.18 (5.10)	-2.67 (6.13)	-2.40 (5.55)

Abbreviations: A, agent's attitude toward the bad side effect (–, negative; 0, not mentioned; ~, don't care; +, positive); SE, side effect only; SE + IE, side effect and intended effect. Approval: How bad did the agent regard the side effect?

3.2.3 Intention

Ascriptions of agent intention also increased somewhat compared to Experiment 1, but again, most participants did not attribute the intention to produce the side-effect to the agent (overall 31.7%, Table 2). The experimental manipulations had no significant effects (p values ≥ 0.278).

3.3 *Discussion*

Experiment 2 investigated intentionality judgments of the bad side-effect in Knobe's (2003) original CEO scenario, while manipulating the agents' attitude towards the side-effect and the moral quality of his primary goal. In contrast to the findings obtained for the scenarios used in Experiment 1, bringing about the side-effect in the CEO scenario was judged as intentional by the clear majority of the participants, and there was a strong discrepancy between judgments of intentionality and of intention. This replicates, for the CEO scenario, two main findings of previous SEE research.

The experimental manipulations had no significant effects on the judgments of intentionality and intent. This means that the lack of effect of agent attitude already found in Experiment 1 was replicated for the CEO scenario, and that one possible explanation of the findings of Experiment 1 — that a

morally good primary goal reduces ascriptions of intentionality to bad side-effects — found no support.

Taken together, Experiments 1 and 2 thus found a high frequency of intentionality ascriptions for the negative side-effect in the CEO scenario, but low frequencies in the minister of education and the sniper scenario. Hence, clear evidence for the SEE was only obtained for the CEO scenario. Furthermore, the attributions of intentionality for the side-effect were not significantly influenced by manipulations of the agent's attitude towards the side-effect (Experiments 1 and 2), the description of the effects of his action (Experiment 1), and the moral quality of the agent's primary goal (Experiment 2). Hence, the attribution of intentionality to bad side-effects turned out to be resistant to the manipulation of several social-cognitive factors that could have been expected to influence it, while at the same time being sensitive to differences in the specific contents of the scenarios (the agents, their social role, the kinds of actions and action effects at issue etc). This suggested to us that, at least for German speakers, the SEE is a context-dependent effect, with high intentionality judgments obtained (perhaps only) in contexts that sufficiently resemble to the CEO case. In Experiments 3 and 4, this hypothesis was investigated further. In Experiment 3, we tested whether an SEE is also shown in other scenarios used in previous studies, whereas in Experiment 4, we tested if the SEE is found in a newly constructed scenario designed to be very similar to the CEO case.

4 Experiment 3

Experiment 3 included six additional SEE scenarios used in previous research, most of which had yielded ascriptions of intentionality for bad side-effects for the great majority of the participants. We reasoned that, if the SEE could be replicated for German speakers in one or more of these scenarios, we could then try to isolate the common features of the scenarios responsible for the effect.

In addition, Experiment 3 tested the effects of yet another experimental manipulation of scenario content: Each scenario was compared to a variant in which the original side-effect was either described as being instrumental to the agent's primary goal, or as a second goal independently pursued by the agent. We expected that in these modified, "non-SEE" scenarios, the gap between intentionality and intention judgments found in the standard SEE scenarios would be closed, or at least strongly reduced. Different from Experi-

ments 1 and 2, Experiment 3 was conducted in the laboratory and had a complete within-subjects design. The experiment was realized using the experiment generator software DMDX (Forster and Forster, 2003). To gain additional insight into the judgment process, we also measured the reaction times of the judgments (e.g., Malle and Holbrook, 2012; see also Reisenzein et al., 1992; Siemer and Reisenzein, 2007). Because these data address different questions from those of main interest in the present article, they will be reported elsewhere.

4.1 *Method*

4.1.1 Participants

Participants were 45 students (34 female), aged between 19 and 36 years (M = 24.3, SD = 4.3). Care was taken to make sure that they had not participated in any of the two previous experiments and were unfamiliar with the SEE effect. The participants were compensated with course credit and 2 Euros.

4.1.2 Design and Scenarios

A 9 (scenario) × (SEE versus Non-SEE) within-subjects design was used. That is, all participants judged both the SEE and the Non-SEE versions of the nine scenarios.

The scenarios comprised the minister and sniper scenario from Experiment 1 and the CEO scenario from Experiment 2. The remaining six scenarios were compiled from previous research: (1) the terrorist scenario (Mallon, 2008:251; 92% intentionality ascriptions of the bad side-effect), (2) the hunter scenario (Nadelhoffer, 2006:142, case C1; 68%), (3) the sales scenario (Phelan and Sarkissian, 2008:294; 75%), (4) the doctor scenario (Uttich and Lombrozo, 2010:95, norm-violating version; percentages of intentionality ascriptions were not reported), (5) the city planner scenario (Phelan and Sarkissian, 2008:296, but with the utterance "I feel terrible" replaced by "I don't care"), and (6) the smoothie scenario (Sloman et al., 2012:156; 95%). The scenarios were translated into German and slightly adapted to accommodate the cultural context of German speakers: actor names were replaced by common German names, and the term "smoothie" in the smoothie scenario was replaced by "lemonade".

For each SEE scenario, we created a modified version ("Non-SEE") that portrayed the original side-effect as either instrumental for the agent's primary goal, or as wanted in itself. To illustrate, the Non-SEE version of the minister of education scenario (Experiment 1) read as follows (the non-SEE part is italicised):

> The Minister of Education of a state is visited by his State Secretary with an important issue. "We have developed a new educational program for our universities. It will markedly increase the quality of study and research. However, many applicants will no longer get a place to study."
> The minister thinks: "That suits me well. Our goals are to increase the quality of study and research in this country, *as well as to prevent the number of students from rising still further*. Let's start the program."
> The program was started. Sure enough, many applicants got no place to study.

The dependent measures consisted of the dichotomous questions asking for intentionality and intention already used in Experiments 1 and 2.

Each scenario was followed by eight dichotomous questions; however, in this article we restrict attention to the two items measuring intentionality and intention. The order of the 18 scenarios, as well as the order of the items within each scenario, was randomized separately for each participant. This ensured that the results would not be systematically biased by prior experiences with a particular scenario, or by previously given answers to the current scenario.

4.1.3 Procedure

The experiment was run in a computer lab of the Institute of Psychology. Participants were run in small groups comprising up to four subjects. To prevent distractions, the work places were shielded by room dividers and the participants wore headphones. The experiment started with a practice scenario intended to familiarize participants with the questions and the answer mode. At the end of the experiment, the participants were debriefed and received the advertised reward.

4.2 *Results*

4.2.1 Intentionality

Figure 1 shows the percentage of intentionality judgments in the SEE and Non-SEE conditions of the nine scenarios. As can be seen, there was a large difference between the two experimental conditions: Whereas overall only 32% of the participants answered the intentionality question with "yes" in the SEE condition, 81% did so in the Non-SEE condition. This difference was highly significant, McNemar test for the equality of dependent proportions, $\chi^2(N = 405, \mathrm{df} = 1) = 163.06$, $p < 0.001$.

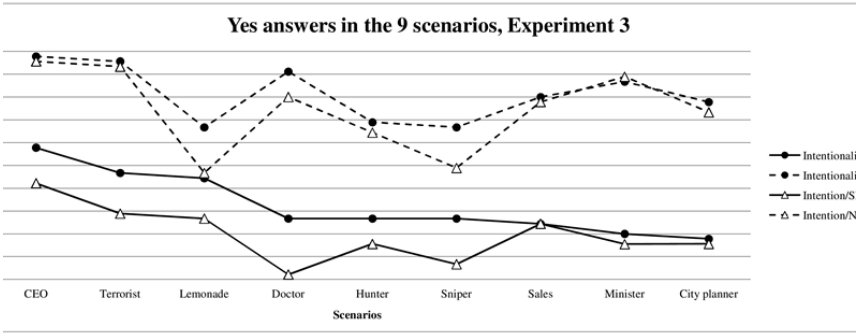

FIGURE 1 *Percentages of Yes answers for the judgments of intentionality and intention in the scenarios used in Experiment 3. Scenarios are ordered by the percentage of intentionality judgments in the SEE condition.*

A closer look at the SEE condition (Fig. 1) reveals that — contrary to what we had hoped — again only the CEO scenario received more than 50% endorsements of intentionality (57.8%). Furthermore, only two other scenarios came close to 50% (the terrorist and lemonade scenarios, 46.7 and 44.4%, respectively). Hence, in six of the nine cases, the bad side-effect was not judged as intentional by the majority of the participants, in contrast to the findings of the original studies in which these scenarios had been used. However, the differences between the three scenarios already used in Experiments 1 and 2 remained significant: minister versus CEO: McNemar $\chi^2(N = 45, \text{df} = 1) = 12.19$, $p < 0.001$; sniper versus CEO: McNemar $\chi^2(N = 45, \text{df} = 1) = 8.45$, $p < 0.01$.

4.2.2 Intention

As in previous SEE studies, including our Experiments 1 and 2, the percentage of intention judgments for the unintended negative side-effect was, with one exception (the sales scenario), lower than the percentage of intentionality judgments (Fig. 1), although the differences were only significant (McNemar test, $p < 0.05$) for four of the nine scenarios (terrorist, smoothie/lemonade, doctor, sniper). Furthermore, as expected, the experimental manipulation had a strong effect on the intention judgments: Whereas the negative action effect was rarely judged as intended in the SEE scenarios (overall 18.4%), it was predominantly judged as intended (overall 76%) in the Non-SEE scenarios, McNemar $\chi^2(N = 405, \text{df} = 1) = 185.44$, $p < 0.001$.

Finally, as can be seen in Fig. 1, the experimental manipulation succeeded in reducing and, in the majority of cases closing, the gap between the judgments of intentionality and intention.

4.3 *Discussion*

Experiment 3 yielded three main findings. First, the results obtained for the scenarios used in Experiments 1 and 2 were essentially replicated with a new sample. Second, the results obtained for these and six additional scenarios, most of which had elicited high intentionality attributions in the original studies, further support the conclusion of Experiment 1 that the SEE is difficult to obtain with German samples: The only clear case of an SEE (intentionality judgments by more than 50%; see above) was again obtained in the CEO scenario. Third, as predicted, the manipulation of the agent's goals regarding the negative effect of his action (it is an unintended side-effect versus instrumental for the agent's primary goal or wanted in itself) had a strong effect not just on judgments of intention, but also on judgments of intentionality.

5 Experiment 4

Experiment 4 had two main goals. The first goal was to provide a further test of the hypothesis that the SEE is context (scenario) dependent. Because Experiment 3 had not revealed any other scenario apart from the CEO scenario in clear evidence for an SEE was obtained, we decided to construct a new scenario that was a close conceptual replication of the CEO scenario. Based on the hypothesis that intentionality ascriptions for bad side-effects depend on the specific contents of the scenario (e.g., the specific actor, his or her social roles), we predicted that a majority of intentionality ascriptions would also be obtained in a scenario similar to the CEO case, even though differing in the details.

Experiment 4 also included, for each SEE scenario, a comparison condition with a morally good side-effect, as in the originally study by Knobe (2003) (e.g., helping the environment in the CEO scenario). In Experiments 1–3 we omitted the positive side-effect condition because intentionality ascriptions in this condition have been found to be very low, and we expected this to be also the case for German participants. Apart from providing a baseline of positive-negative differences to compare to the means scenarios (see below), the positive side-effect conditions allowed us to test whether the percentages of intentionality ascriptions for negative side-effects obtained in Experiment 3, even if they were below 50%, were at least significantly higher than in the positive side-effect condition, and thus whether there was at least evidence for a weak form of the SEE (see above).

The second goal of Experiment 4 was to test whether the difficulties of finding an SEE for German speakers extends to similar effects in related domains.

To investigate this issue, we chose the "means effect" recently reported by Cova and Naar (2012). The means effect is an SEE-type effect for judgments of intentionality of the means (e.g., concrete action) chosen to achieve a goal. According to the traditional model of intentionality ascriptions, intentionality judgments for means should be high regardless of their moral quality, because means are deliberately chosen. However, Cova and Naar (2012) found that although intentionality ascriptions for morally bad means were high, those of morally good means were low (overall difference 82 vs. 42%; note that the unexpected finding here concerns the low intentionality ascriptions for good means). In Experiment 4, we attempted to replicate this generalization of the SEE in a German sample.

5.1 *Method*

5.1.1 Participants

77 first-year students (57 female), with a mean age of 23.7 years (SD = 8.4), participated in the study in exchange for course credit. None of the participants had taken part in the preceding experiments or was familiar with the SEE or means effect.

5.1.2 Design and Procedure

Participants judged nine scenarios in two blocks, comprising the SEE and means scenarios, respectively. Each scenario comprised two conditions to which the participants were randomly assigned: a positive condition (featuring a morally good side-effect or means) and a negative condition (featuring a bad side-effect or means). The order of scenarios within each block was randomized to prevent systematic transfer effects.

The SEE scenarios comprised the CEO scenario and the terrorist scenario, the two scenarios for which the highest percentages of intentionality ascriptions had been obtained in Experiments 1–3, plus a new scenario named "pharmaceutical company". This scenario was constructed to resemble the CEO scenario as closely as possible without duplicating it: A head of a pharmaceutical company decides to raise the price of a drug manufactured exclusively by this company to increase profits, accepting a bad side-effect (some people cannot afford the drug any longer) that he clearly foresees, but does not care about. Based on the hypothesis of the context-dependence of the SEE, we predicted that the side-effect in the pharma scenario would, as in the CEO case, be judged as intentional by the majority of the participants.

The means scenarios comprised German translations of the three means scenarios used by Cova and Naar (2012) (Experiment 1) (inheritance, coffee

break, prisoner) for which the authors had found significant differences between the bad and good versions, plus three new scenarios (featuring a farmer, a lawyer, and a party leader) modeled after these. All means scenarios had the same basic structure: An agent aims to achieve a goal by performing either a morally good or bad action (the means). For example, the farmer scenario was as follows:

> [Bad means]. The farmer's wife says to the farmer: "I know a way to increase our income. We could grow genetically modified corn. This will reduce our crop losses." The farmer responds: "I don't care whether or not the corn is genetically modified. I only want to increase our income. Let's do it." He grows genetically modified corn and the income increases.
> [Good means]. The farmer's wife says to the farmer: "I know a way to increase our income. We could grow this old and robust breed of potatoes that is threatened by extinction. This will reduce our crop losses." The farmer responds: "I don't care whether or not the potato breed is threatened by extinction. I only want to increase our income. Let's do it." He grows the old breed of potatoes and the income increases.

For each scenario, the participants answered two dichotomous yes/no-questions. First, as a manipulation check, participants were asked whether the agent could be praised (in the morally good condition) or blamed (in the morally bad condition) for producing the side-effect or choosing the means, respectively (see Cova and Naar, 2012). The second question asked whether the agent had intentionally produced the side-effect or the means, respectively.

5.2 *Results*

5.2.1 Praise/Blame

The manipulation check replicated the results of Knobe (2003) and Cova and Naar (2012): bad side-effects and means were regarded as blameworthy by the vast majority of the participants (on average across scenarios, 94%), whereas their good counterparts were judged as praiseworthy only by a minority (29%). The good-bad differences were significant for each scenario (Chi-square test, p values < 0.001).

5.2.2 Intentionality

Figure 2 shows the percentages of intentionality judgments in the two moral conditions (good vs. bad side effect/means) for the nine scenarios together

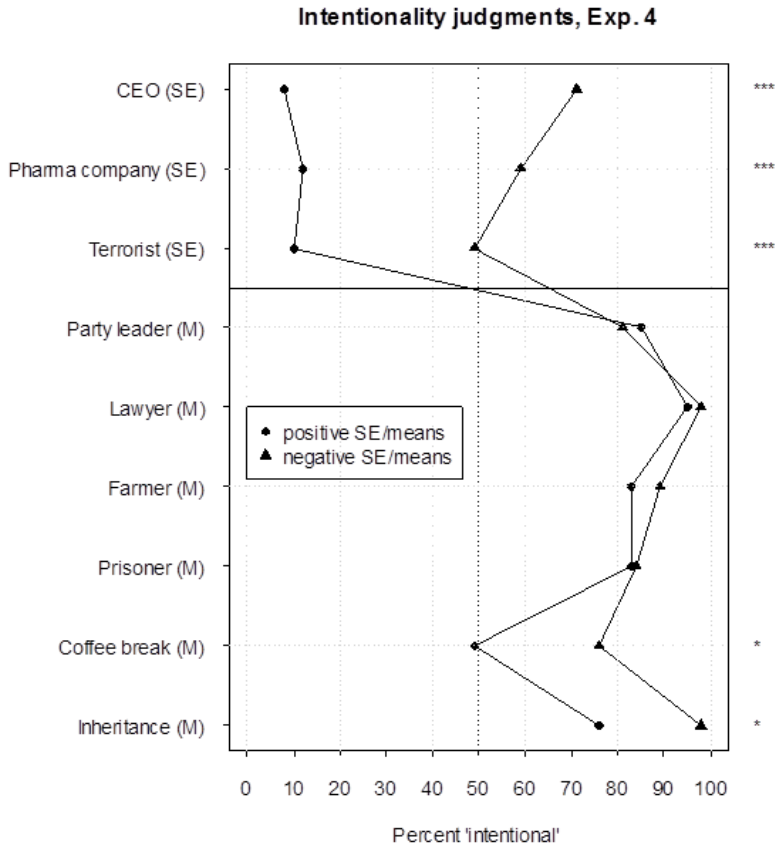

Percentages of intentionality judgments in the side-effect and means scenarios, Experiment 4. SE = SEE scenario; M = means-scenario. $^{*}p < 0.05$, $^{***}p < 0.001$.

with the significance levels of the differences. In the bad side-effect conditions, the side-effect was judged as intentional by 49% in the terrorist and by 71% in the CEO scenario, replicating the results of Experiments 2 and 3. In the new "pharmaceutical company" scenario, the negative side-effect was judged as intentional by 59% and thus approximated the CEO scenario, as we had predicted. In the good side-effect conditions of the scenarios, intentionality judgments were infrequent (8–10%) and in all cases significantly ($p < 0.001$) less frequent than in the corresponding negative side-effect condition.

For the means scenarios, we obtained, like Cova and Naar (2012), a high percentage of intentionality ascriptions for bad means. However, different from Cova and Naar (2012), in five of the six scenarios (the exception was the coffee break scenario), we also obtained a clear majority of intentionality ascriptions

for good means (see Fig. 2). Hence, if one defines the means effect by a majority of "intentional" judgments for bad means, coupled with a majority of "unintentional" judgments for good means (analogous to the SEE, see above), no evidence for the effect was found. Even a weak form of the means effect (a higher frequency of intentionality ascriptions for bad than good means) was reliably present ($p < 0.05$) only in two of the six scenarios (inheritance and coffee break).

5.3 *Discussion*

The results obtained for the SEE scenarios confirmed our expectation that German participants, too, rarely attribute intentionality to a good side-effect. This replicates the findings of Knobe (2003) and subsequent authors. Nevertheless, the baseline of 10% intentionality ascriptions for good side-effects constitutes a useful additional piece of information: It allows to test whether the frequencies of intentionality ascriptions for bad side-effects obtained in Experiment 3, even though they were mostly below 50%, are at least significantly higher than those for good side-effects. It turned out that, given the sample sizes used, no statistically reliable ($p < 0.05$, Chi-square test) good-bad difference could be established for the doctor, hunter, sniper, sales, minister, and city planner scenario from Experiment 3.

As in Experiment 3, a substantial difference in the size of the SEE effect was found between the terrorist scenario (49% intentionality judgments) and the CEO scenario (71%). The most interesting finding, however, was obtained for the pharmaceutical company scenario. Based on the context-dependency hypothesis, we had predicted that this scenario, which was designed to closely resemble the CEO case, would result in a majority of intentionality attributions. This prediction was supported. Hence, we were finally able to replicate the SEE in a different situation. This finding supports our hypothesis that the specific content of the scenarios is crucial for the occurrence of intentionality ascriptions to bad side-effects.

The second goal of Experiment 4 was to attempt replicating the "means effect" reported by Cova and Naar (2012). Although the manipulation check confirmed that the moral quality of the good and bad means was perceived as intended, the means effect at best be partially replicated: Different from Cova and Naar (2012), not only bad, but also good means were judged as intentional by the majority of the participants (overall 87%); and for four of the six scenarios, not even a statistically significant difference between good and bad means was obtained. Hence, consistent with our results concerning the SEE, the means effect was absent, or at least much reduced, in German participants, and seemed to depend on the scenario.

6 General Discussion

The SEE is an intriguing phenomenon that has generated a considerable amount of research. Several explanations of the SEE have been proposed. Whereas most of these trace the effect to the peculiarities of folk-psychological social and moral reasoning (e.g., Knobe, 2006; Uttich and Lombrozo, 2010; Holton, 2010), some authors have argued that the SEE may be caused by linguistic (Cova et al., 2012) or methodical factors (Guglielmo and Malle, 2010, studies 5 and 6). These explanations suggest to focus attention on the details of the scenarios used in SEE research. The present results support this suggestion.

In four experiments including 565 German speakers and 11 different SEE scenarios, clear evidence for the SEE (i.e., intentionality ascriptions for bad side-effects by the majority of the participants) was obtained in only two cases: Knobe's (2003) original CEO scenario, and a scenario designed to be very similar in content to the CEO case (the pharmaceutical company scenario). In the remaining nine scenarios, the percentage of intentionality judgments ranged from slightly below 50% (terrorist, lemonade) to 18% (city planner; see Fig. 1). Hence, in nearly all scenarios, the intentionality ascriptions for bad side-effects of our participants were much more "rational", or better in line with the "simple view" of intentionality judgments, than would have been predicted by prior research. In addition, Experiment 4 found that the difficulties of obtaining an SEE effect for German speakers generalized to intentionality judgments of means: The majority of our participants judged not only bad means, but also good means as intentional, and in four of the six scenarios, the difference in the frequency of intentionality ascriptions for good and bad means did not even reach conventional levels of statistical significance.

Paralleling the difficulties of obtaining an SEE for the majority of German speakers, Experiments 1 and 2 found the SEE to be largely resistant to experimental manipulations of several seemingly plausible social-cognitive determinants of the effect: the agent's attitude towards the side-effect, the moral quality of the agent's primary goal, and the description of the effects of the action. The manipulation of the moral quality of the side-effect (Experiment 4) did have a significant effect on the ascriptions of intentionality for side-effects in the three scenarios used in this experiment, but not in six other scenarios used in Experiment 3 (assuming a baseline of 10% intentionality ascriptions for the corresponding good side-effects). The only experimental manipulation that yielded a consistent, scenario-transcendent effect was a change of the description of the side-effect from unintended (or ambiguous), to instrumental or

wanted in itself. As predicted by the traditional model of intentionality judgments, this manipulation led to large changes in attributed intention, that were mirrored in similarly large changes in intentionality ascriptions, which in the majority of cases closed the gap between these two kinds of judgments. This finding supports a weakened version of the "simple view" of intentionality judgments (Bratman, 1984): Even though the simple view of intentionality judgments as originally stated appears to be wrong, the perceived presence of intent clearly remained a strong determinant of intentionality attributions in our data.

Finally, Experiment 4 provides initial support for our hypothesis that at least for German speakers, the SEE (as well as, apparently, the "means effect") is context-dependent, that is, it depends on the specifics of the scenario (or perhaps on an interaction between scenario and language): Only in the CEO scenario and a structurally highly similar scenario (pharmaceutical company) was the bad side-effect judged as intentional by the majority; in most other scenarios, a clear majority even judged the bad side-effect as unintentional. Although the present studies do not allow to unambiguously isolate the scenario properties that drive the obtained between-scenario differences, our guess is that they concern the stereotypic beliefs about the agents depicted in the scenarios, especially their social role (e.g., CEO, head of a pharmaceutical company, minister, sniper) and the differences in power, moral obligations, etc., that go with it. These beliefs could be automatically activated by the scenario descriptions and then guide intentionality judgments (see Siemer and Reisenzein, 2007, for related considerations concerning the case of emotion inference from scenarios) — apparently to a degree that information opposing the stereotypical beliefs provided in the scenarios is ignored (Experiments 1 and 2). Preliminary evidence for this hypothesis is provided by the finding of Experiment 4 that a majority of intentionality ascriptions could be obtained in a scenario describing an agent holding a social role, and making decisions, very similar to the CEO case. Thus, our suggestion is that what is decisive for judgments of intentionality about unintended bad side effects is not the negative side-effect per se, but the side-effect in interaction with specific descriptive and evaluative schemas evoked by the a scenario, in particular concerning the social role of the agent (see also Sripada and Konrath, 2011).

References

Bratman, M. (1984). Two faces of intention. *Philosophical Review* 93, 375–405.

Cova, F., Dupoux, E. and Jacob, P. (2012). On doing things intentionally. *Mind and Language* 27, 378–409.

Cova, F. and Naar, H. (2012). Side-Effect effect without side effects: The pervasive impact of moral considerations on judgments of intentionality. *Philosophical Psychology* 25, 837–854.

Fincham, F.D. and Jaspers, J.M. (1980). Attribution of responsibility: From man the scientist to man as lawyer. In Berkowitz, L. (Ed.), *Advances in Experimental Social Psychology 13*, pp. 81–138. Academic Press, New York, NY.

Forster, K.I. and Forster, J.C. (2003). DMDX: A Windows display program with millisecond accuracy. *Behavior Research Methods, Instruments and Computers* 35, 116–124.

Guglielmo, S. and Malle, B.F. (2010). Can unintended side effects be intentional? Resolving a controversy over intentionality and morality. *Personality and Social Psychology Bulletin* 36, 1635–1647.

Heider, F. (1958). *The psychology of interpersonal relations*. Wiley, Hoboken, NJ.

Holton, R. (2010). Norms and the Knobe effect. *Analysis* 70, 417–424.

Knobe, J. (2003). Intentional action and side-effects in ordinary language. *Analysis* 63, 190–193.

Knobe, J. (2006): The concept of intentional action: a case study in the uses of folk psychology. *Philosophical Studies* 130, 203–31.

Knobe, J. and Burra, A. (2006). The folk concepts of intention and intentional action: A cross-cultural study. *Journal of Cognition and Culture* 6, 113–132.

Knobe, J. and Mendlow, G.S. (2004). The good, the bad and the blameworthy: Understanding the role of evaluative reasoning in folk psychology. *Journal of Theoretical and Philosophical Psychology* 24, 252–258.

Leslie, A.M., Knobe, J. and Cohen, A. (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 and Language* 23, 165–189.

Malle, B.F. (2006). Intentionality, morality, and their relationship in human judgment. *Journal of Cognition and Culture* 6, 87–112.

Malle, B.F. and Holbrook, J. (2012). Is there a hierarchy of social inferences? The likelihood and speed of inferring intentionality, mind, and personality. *Journal of Personality and Social Psychology* 102, 661–684.

Malle, B.F. and Knobe, J. (1997). The folk concept of intentionality. *Journal of Experimental Social Psychology* 33, 101–121.

Mallon, R. (2008). Knobe versus Machery: Testing the trade-off hypothesis. *Mind and Language* 23, 247–255.

McCann, H.J. (2005). Intentional action and intending: Recent empirical studies. *Philosophical Psychology* 18, 737–748.

McGraw, K.M. (1987). Guilt following transgression: An attribution of responsibility approach. *Journal of Personality and Social Psychology* 53, 247–256.

Meinong, A. (1894). *Psychologisch-ethische Untersuchungen zur Wert-Theorie* [Psychological-ethical investigations in value theory]. Leuschner and Lubensky, Graz. Reprinted in R. Haller and R. Kindinger (Eds.). (1968). *Alexius Meinong Gesamtausgabe* [Alexius Meinong's complete works], Volime III, pp. 3–244. Akademische Druck- und Verlagsanstalt, Graz.

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 juror impartiality. *Philosophical Explorations* 9, 203–219.

Nadelhoffer, T. and Nahmias, E. (2007). The past and future of experimental philosophy. *Philosophical Explorations* 10, 123–149.

Nichols, S. and Ulatowski, J. (2007). Intuitions and individual differences: the Knobe Effect revisited. *Mind and Language* 22, 346–365.

Pellizzoni, S., Girotto, V. and Surian, L. (2009). Beliefs and moral valence affect intentionality attributions: the case of side effects. *Review of Philosophy and Psychology* 1, 201–209.

Phelan, M.T. and Sarkissian, H. (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.

Reips, U.-D. and Neuhaus, C. (2002). WEXTOR: A web-based tool for generating and visualizing experimental designs and procedures. *Behavior Research Methods, Instruments and Computers* 34, 234–240.

Reisenzein, R. (1995). On Oatley and Johnson-Laird's theory of emotion and hierarchical structures in the affective lexicon. *Cognition and Emotion* 9, 383–416.

Reisenzein, R., Debler, W. and Siemer, M. (1992). Der Verstehensvorgang bei scheinbar paradoxen Wirkungen von Lob und Tadel. *Zeitschrift für experimentelle und angewandte Psychologie* 39, 129–150.

Siemer, M. and Reisenzein, R. (2007). The process of emotion inference. *Emotion* 7, 1–20.

Skitka, L.J., Bauman, C.W. and Sargis, E.G. (2005). Moral conviction: Another contributor to attitude strength or something more? *Journal of Personality and Social Psychology* 88, 895–917.

Sloman, S.A., Fernbach, P.M. and Ewing, S. (2012). A causal model of intentionality judgment. *Mind and Language* 27, 154–180.

Sripada, C.S. and Konrath, S. (2011). Telling more than we can know about intentional action. *Mind and Language* 26, 353–380.

Uttich, K. and Lombrozo, T. (2010). Norms inform mental state ascriptions: A rational explanation for the side-effect effect. *Cognition* 116, 87–100.

Weiner, B. (1985). An attributional theory of achievement motivation and emotion.

Psychological Review 92, 548–573.

Weiner, B. (1995). *Judgments of responsibility: A foundation for a theory of social conduct*. Guilford Press New York, NY.