

Perspective and Epistemic State Ascriptions

Markus Kneer¹ 

© Springer Science+Business Media B.V. 2017

Abstract This article explores whether perspective taking has an impact on the ascription of epistemic states. To do so, a new method is introduced which incites participants to imagine themselves in the position of the protagonist of a short vignette and to judge from her perspective. In a series of experiments (total $N=1980$), perspective proves to have a significant impact on belief ascriptions, but not on knowledge ascriptions. For belief, perspective is further found to moderate the epistemic side-effect effect significantly. It is hypothesized that the surprising findings are driven by the special epistemic authority we enjoy in assessing our own belief states, which does not extend to the assessment of our own knowledge states.

1 Introduction

1.1 Perspective and Perspective Taking

The impact of perspective is ubiquitous in psychology. We tend to overestimate our own abilities and positive contributions in comparison to those of others, as documented by the self-serving bias (for reviews cf. Blaine and Crocker 1993, Mezulis et al. 2004). In explaining other people's actions, we primarily invoke the agent's character traits, whereas in explaining our own actions, we focus strongly on situational features, a phenomenon which is called the actor-observer asymmetry (Jones and Nisbett 1971, cf. Knobe and Malle (2002) for a review and Malle (2006) for a meta-analysis). Perspective also influences judgments about the riskiness of decisions (Fernandez-Duque and Wifall 2007; Pollai and Kirchler 2012) and the moral permissibility of actions (Nadelhoffer and Feltz 2008). Even as regards happiness maximization, perspective has a pronounced effect: While people are strongly inclined to plug others into a Nozickian experience machine, they are loathe to plug in themselves (Weijers 2014).

✉ Markus Kneer
markus.kneer@gmail.com

¹ Department of History & Philosophy of Science, University of Pittsburgh,
1017 Cathedral of Learning, Pittsburgh, PA 15260, USA

Related effects can be observed as concerns perspective-taking, or ‘sympathizing’ with others in the sense of Adam Smith (cf. Smith 1790/2002 especially I.I.III-IV and II.II.II), whereby an observer imagines herself in the situation of a particular target individual. Assuming the perspective of other people correlates with an increased appreciation of them (Batson et al. 2007), even if they are members of outgroups (Galinsky and Moskowitz 2000; Galinsky et al. 2005), more cooperation (Parker and Axtell 2001), more successful resolution of conflicts (Galinsky et al. 2008; Takaku et al. 2001), a higher willingness to forgive transgressions (McCullough et al. 1997) and a higher propensity to help them (Cialdini et al. 1997; Maner et al. 2002). Perspective taking has also been found to reduce egocentric biases (Epley et al. 2006) and stereotyping (Galinsky and Moskowitz 2000).

Explanations of perspective taking effects standardly invoke an altered perception of the target individual due to a projection of the perceiver’s self onto the target (schematically: self→target). Interestingly, the relation might also hold in the reverse direction, in so far as properties of the target are projected onto the perceiver’s self (schematically: target→self). Laurent and Myers (2011), for instance, report data according to which imagining oneself in the shoes of a researcher leads to a change in self-concept and correlates with an increased self-ascription of attitudes and characteristics of researchers.

1.2 Perspective and Mental State Ascriptions

A topic that has attracted comparatively little attention regards whether the ascription of mental states – intention, desire, knowledge and belief – are influenced by perspective and perspective taking. In this paper, we will focus on epistemic mental states, that is, knowledge and belief. An inquiry of this sort can be motivated by two key reasons: One is theoretical, and arises from the special epistemic status characteristic of ‘self-knowledge’, that is, knowledge of one’s own mental states. The second reason is driven by more practical concerns and regards the role knowledge and belief play in moral psychology and criminal jurisprudence in virtue of their inculpatory nature. We’ll briefly consider them in turn.

First, self-knowledge is standardly taken to enjoy a unique epistemic status in theory of mind (for discussion, cf. inter alia Dretske 1994, Byrne 2005, 2011, Cassam 1997, 2011, 2014, Gertler 2010, Siewert 2012). While we have to infer the mental states of others from their assertions and behavior, we stand in a distinctly agential relation to our own mental states and they tend to be directly transparent to us. Privileged access of this sort is considered to imbue self-knowledge with a special epistemic authority intimately tied to the first-person perspective. An interesting question thus arises whether, in taking the perspective of a target subject, we ascribe epistemic states in different fashion than from a straightforward observer perspective. If so, this would suggest that the attribution of belief and knowledge is not only susceptible to the *evidence* at hand (which can be held constant across observer and quasi-actor perspectives), but also to the *mode* (first-personal v. third-personal) from which such ascriptions are effected.

Second, knowledge and belief, just like intention, play an important role in moral psychology and criminal jurisprudence. Moral culpability or responsibility standardly require not only a harmful outcome *h*, but also an inculpatory state of mind such as an *intention* to bring about *h*, or the *knowledge* or *belief* that one will bring about *h*. The situation is similar as regards the law: Except in cases of strict liability, a defendant is

deemed legally responsible for a harmful outcome only if it can be established that she committed a guilty act (*actus reus*) with an inculpatory state of mind (*mens rea*).¹

In the assessment of moral culpability, actors and observers frequently disagree about whether a harmful action was indeed intended, foreseen or could have been anticipated by the actor. To bridge the gap between diverging actor- and observer assessments of guilty states of mind, we oftentimes invite the other party to evaluate the situation from our own point of view. For intentionality, putting oneself in the shoes of another can have a significant impact: In a recent study, Feltz et al. (2012) report that the Knobe effect, according to which knowingly incurred negative side-effects are judged more intentional than positive ones, drastically decreases in size for self-ascriptions in contrast to other-ascriptions. Differently put, the impact of outcome valence, which influences ascriptions of the *mens rea* of intentionality is moderated by perspective taking. Recent studies suggest that the Knobe effect also arises for knowledge and belief ascriptions (the ‘epistemic side-effect effect’); we will thus follow Feltz and colleagues in adopting a 2 (perspective: actor v. observer) × 2 (outcome: good v. bad) design for the experiments. The studies presented below will generate insights across three contrastive dimensions: (i) first- v. third-person perspective, (ii) knowledge v. belief, and (iii) positive v. negative outcomes.

The paper is structured as follows: Section 2 introduces the Knobe effect and the epistemic side-effect effect, as well as a recent theory that gives a unified explanation for both of them. Various hypotheses regarding the impact of perspective on epistemic state ascriptions are discussed with reference to the theoretical literature and recent empirical findings. Section 3 outlines a new experimental method which incites participants to imagine themselves in the shoes of the protagonist of an imaginary scenario. Five experiments (total $N = 1980$ participants) using three different scenarios from the epistemic side-effect effect (ESEE) literature are reported in Sections 3 to 7. Perspective turns out to have a robust effect on belief ascriptions, and significantly moderates the ESEE, though no evidence in favour of a perspective effect for knowledge ascriptions could be found. Section 8 discusses the findings and points out interesting future research possibilities.

2 Setting the Stage

2.1 The Knobe Effect

Side-effect scenarios, and the side-effect effect have been at the centre of much debate in experimental philosophy. The side-effect effect, or ‘Knobe effect’, refers to the asymmetry in intentionality ascriptions depending on whether a side-effect is positive or negative. To illustrate the point, let’s consider Joshua Knobe’s (2003a) original vignette, which comes in two variations (harm v. help, alternative formulation in square brackets):

The vice-president of the company went to the chairman of the board and said,
“we are thinking of starting a new program. It will help us increase profits, [but/

¹ For the similarities and differences of guilty mind ascriptions in moral judgment and criminal jurisprudence, cf. Duff (1989) and Moore (2011).

and] it will also [harm/help] the environment.” The chairman of the board answered, “I don’t care at all about [harming/helping] the environment. I just want to make as much profit as I can. Let’s start the new program. They started the program. Sure enough, the environment was [harmed/helped]. (2003a, p. 191)

Participants were asked whether they thought the chairman had intentionally harmed (helped) the environment. When the side-effect was negative, 82% of the subjects judged that the chairman had brought it about intentionally, when it was positive, only 23% did.² The ‘folk’ hence tends to judge a knowingly incurred negative side-effect as intentionally brought about, whereas positive side-effects are deemed unintentional byproducts of the main action.³

2.2 The Epistemic Side-Effect Effect

The valence of the side-effect has also been shown to influence judgments regarding knowledge ascriptions. In a series of experiments with side-effect cases, Beebe and Buckwalter (2010) and Beebe and Jensen (2012) report that the folk are significantly more willing to ascribe knowledge to the chairman in the above scenario when the side-effect is negative than when it is positive. This is striking: A debate could be had as to whether certain factors such as moral valence, norms or the blameworthiness of the agent might be relevant for the ascription of intentionality – in fact, the ‘morally charged’ accounts of intentional action of Harman (1976) and Duff (1980, 1982) have long predated the discovery of Knobe Effect. By contrast, moral or normative valence differ radically from the factors commonly deemed relevant for knowledge ascription in traditional epistemology (cf. e.g. Bonjour 2009, Feldman 2006, Lehrer 2015). So far, the ESEE has proven robust and pervasive. It also arises for knowledge ascriptions in Gettier cases (Beebe and Shea 2013; Buckwalter 2014), the attribution of belief (Beebe 2013), and it has been replicated in various studies (cf. Dalbauer and Hergovich 2013, Turri 2014 and Colaço et al. (unpublished); for discussion cf. Buckwalter (2012), Turri (2016), Beebe (2016).

2.3 The Epistemic Side-Effect Effect & Doxastic Heuristics

The ESEE is a relatively recent discovery, and theories that account for it are still somewhat scarce. We will principally focus on a proposal by Alfano et al. (2012) and

² The effect has been widely replicated (Knobe 2003a, b, 2004b; Mele and Cushman 2007). It holds across cultures (Knobe and Burra 2006) and ages (Leslie et al. 2006), and is just as robust for professional lawyers and judges as for laymen (Kneer and Bourgeois-Gironde 2017a, forthcoming, 2017b). For a brief overview, see Feltz (2007) as well as Cova (2016). Cova and Naar (2012) report an asymmetry in intentionality ascriptions without relying on side-effect scenarios.

³ Whether the side-effects in Knobe’s scenarios should indeed be understood as *knowingly* incurred is by now a topic of controversy, because a similar asymmetry arises in epistemic state ascriptions for the same scenarios (cf. Beebe and Buckwalter 2010; Beebe and Jensen 2012). More on this below.

Robinson et al. (2015), whose heuristic-driven explanation of the ESEE regarding belief has the virtue of generalizing to a broad variety of Knobe effect phenomena. The view is of particular interest for our purposes, as perspective plays a central role in it. It is best introduced in two steps:

In a first step, Alfano and colleagues provide a theory of the ESEE for belief. Dovetailing with other recent norm-driven accounts of the Knobe effect (cf. e.g. Holton 2010, Lombrozo and Uttich 2010, Uttich and Lombrozo 2010), the ESEE regarding *belief* is explained by a ‘norm-violation & belief-attribution heuristic’ (NBA). According to this heuristic, other agents are taken to have good practical reason to reflect about actions and side-effects that violate salient norms and hence to form beliefs about them. In cases where there is no salient norm infraction, by contrast, we do not expect others to reflect in depth about their actions or to form particular outcome-regarding beliefs about potential side-effects. The rationale for postulating a heuristic of this sort, Alfano et al. argue, is that we take the formation of our *own* beliefs to be governed by a similar principle (the ‘norm-violation & belief-formation heuristic’, or NBF):

NBF: If my own φ -ing would make it the case that p and p violates a norm salient to me, believe that φ -ing would make it the case that p . (2012: 7)

The second step provides a unified explanation of the Knobe effect on the basis of the ESEE for belief. Each of the mental states susceptible to the Knobe effect (intentionality, knowledge, desire, etc.) is intimately tied to belief and reflection, and their attribution is governed by doxastic heuristics⁴ of the following sort:

(Know \rightarrow Believe) Agent a knows that φ -ing would make it the case that p only if a believes that φ -ing would help to make it the case that p .

(Intentionally \rightarrow Believe) Agent a intentionally makes it the case that p by φ -ing only if a believes that φ -ing would help to make it the case that p . (2012: 4)

Now, if belief ascription plays an important role with regards to other mental states that manifest the Knobe effect,⁵ the latter can be understood as a consequence of the ESEE regarding belief. We will call this the *constitutive belief asymmetry conjecture*. On this view, the asymmetry regarding, for instance, intentionality ascriptions across differently valenced side-effects, is driven by the asymmetry in belief ascriptions.

In short, Alfano et al. hypothesize that Knobe-type effects with respect to most mental state ascriptions are driven by the ESEE regarding belief. The ESEE for belief is explained by the norm-sensitive heuristic guiding other-ascriptions of belief (NBA), which in turn derives from a similarly norm-sensitive principle we take to govern the formation of our own beliefs (NBF). Experiments focusing on perspective will be helpful in testing the assumed relation between the other-regarding NBA heuristic and the self-regarding NBF heuristic, which is at the centre of this approach.

⁴ Only a few are quoted here, though Alfano and colleagues defend similar heuristics for desiring p , being in favour of p , advocating p and other attitudes that have been shown to be susceptible to the Knobe effect.

⁵ The content of the various doxastic heuristics does not require belief as a necessary condition for the target mental states. A strong, yet defeasible relationship holding for ordinary cases is sufficient. The agents, of course, need not to be explicitly aware of the heuristics.

2.4 Hypotheses

In what ways might perspective influence knowledge and belief ascriptions? Perspective might affect *both* belief and knowledge ascriptions, *neither*, or only *one of the two*. I'll discuss and assess these options in turn.

According to previous findings by Feltz et al. (2012), perspective-taking does have an effect on intentionality ascriptions, and assuming the point of view of a target subject partially crowds out the Knobe effect. On Alfano et al.'s constitutive belief asymmetry conjecture just discussed, the existence and magnitude of the Knobe effect for intentionality is ultimately driven by a corresponding ESEE for belief ascriptions. We can hence infer a similar effect of perspective on belief ascriptions and, if belief is taken to be constitutive of knowledge, the same would hold for knowledge ascriptions.

While certain features of the heuristic account predict that *both* belief and knowledge will be affected by perspective (at least given previous data for intentionality), others would rather predict that *neither* will be. According to the norm-belief formation heuristic, we tend to dwell on and self-ascribe beliefs concerning norm-violating side-effects. In the interest of saving cognitive resources, we do not reflect on positive side-effects, and do not self-ascribe beliefs concerning them. Hence, NBF would predict a strong asymmetry in belief *self*-ascription across differently valenced side-effect scenarios, just as there is a strong asymmetry for *other*-ascriptions. Perspective should have little or no impact, since our asymmetric other-ascriptions of epistemic states (governed by NBA) only generalize to the observer perspective what we consider adequate for belief self-ascriptions (NBF).

Could perspective influence the attributions of only one of the two epistemic states, but not the other? According to orthodox epistemology, knowledge differs from belief in so far as it must be justified and true. A predicted asymmetric impact of perspective would thus invoke the influence of perspective on the ascriptions of truth or epistemic justification. However, the latter play but a marginal role in side-effect scenarios, and it is not immediately evident why, and in what ways, they should be susceptible to the influence of perspective. But if, among the standard factors of knowledge, only belief is hypothesized to be sensitive to perspective, and if knowledge is taken to entail belief, then there is little *prima facie* reason to predict an asymmetric influence of perspective on the two types of epistemic state ascriptions.

Things are different once we turn to nonorthodox conceptions of knowledge. Williamson (2002), for instance, has argued that knowledge is a *sui generis* mental state independent of belief. Recent findings in experimental epistemology have also put pressure on the view that knowledge entails belief (Myers-Schulz and Schwitzgebel 2013, Murray et al. 2013, for critical discussion, cf. Rose and Schaffer 2013, Buckwalter et al. 2015). When presented with a vignette of Radford's unconfident examinee, for instance, participants tend to ascribe knowledge but not belief. If there were indeed no close conceptual relation between knowledge and belief, it would not be surprising to find that knowledge and belief ascriptions are affected differently by perspective taking.

But perhaps traditional epistemology, too, can make room for a split prediction: Whereas one has relatively unfettered epistemic authority as concerns one's own beliefs, this is not the case for knowledge. Whether a particular representational state

does indeed constitute an instance of knowledge cannot be settled with a mere glance at the inventory of one's own mind. It involves an outward-looking gaze to bridge the mind-world gap – the very same process which characterizes the assessment of mental states of *others* – in which a plethora of things can go wrong. Hence, out of the four basic types of epistemic state ascription – perspective (self v. other) x state (belief v. knowledge) – self-ascription of belief is characterized by considerably more epistemic authority than the three other types of ascription (self/knowledge; other/belief; other/knowledge). This privileged epistemic position, intimately linked to the first-person perspective, might support the prediction that only belief ascriptions vary across the assessment of others versus oneself, not knowledge ascriptions.

3 Experiment 1: Belief in the Chairman Scenario

The first experiment explores the impact of perspective on belief ascriptions. One might be tempted simply to change the pronouns of the vignette from 'he, the chairman' to 'you, the chairman' in order to incite participants to assume the chairman's perspective. Feltz et al. (2012) did precisely that in a study that focused on intentionality ascriptions. No significant effect due to perspective was found. In an experiment where subjects had to participate actively in an investment game or else merely judge from an observer's perspective, however, a significant variation across perspectives was detected. From this, the authors conclude that 'actor-observer differences are not likely to be found by simply asking participants to imagine that they are actors', and that bringing out perspectival asymmetries requires 'putting participants in the actual decision-making environment' (2012, p. 682–3). Importantly, Feltz et al. take their results to 'challenge the substantial philosophical and empirical reliance on hypothetical thought examples about intentional action.' (2012, p. 673).

This verdict might be too quick. On the one hand, it conflates the distinction between judging from one's *own* perspective and perspective *taking*, in which, by definition, we imagine the perspective of another individual (rather than judge from our own perspective). As discussed above, the former cognitive mode is no substitute for the latter; it is a different phenomenon altogether. On the other hand, a mere change in pronouns might indeed be insufficient in order to trigger perspective taking, but the options are not thereby exhausted. For the present study, participants were asked to take a few 'managerial decisions' from the chairman's point of view, a process which was intended to help them into his proverbial shoes. Participants presented with the first-person scenarios thus underwent a *point of view augmentation procedure* (POVAP) before they responded to the main question (further detailed below).

3.1 Participants

Three hundred fifty one participants were recruited on Amazon Mechanical Turk to complete a paid Qualtrics online survey. Each subject was randomly assigned to one of the four conditions (2 outcomes: help v. harm × 2 perspectives: actor v. observer) of the CHAIRMAN scenario. Datasets of non-native speakers, subjects who responded in under twenty seconds, and participants failing an attention test or changing their answer ten times

or more were excluded. 254 subjects remained, of which 113 were female. The average age of the participants was 36,8 years (SD: 11.9 years).

3.2 Materials and Procedure

Following Beebe and Buckwalter (2010) and Beebe and Jensen (2012), an adapted version of Knobe's (2003a) CHAIRMAN scenario was used for the experiment. Alongside the customary variation in outcome (help v. harm), the vignettes differed with respect to perspective (actor v. observer). That is to say, participants either had to imagine *themselves* in the position of the company's president or else judge the epistemic states of *another person*, chairman Smith. Participants in the observer conditions read the following scenario (outcome variations in square brackets):

Imagine that Mr Smith is the chairman of a small company, which produces and sells film posters. Smith has 50 employees, including Frank, an assistant who advises Smith on the business strategy.

One day Smith's assistant Frank comes into his office and says:

"There is a new printing technique which we could use. It is much cheaper and will increase profits, [but/and] it will also [harm/help] the environment."

Smith thinks to himself: "I don't care at all about [harming/helping] the environment. I just want to make as much profit as I can, and my calculations show that the new technique will increase profits."

Smith decides to implement the new printing technique. As predicted, profits increase and the environment is [harmed/helped].

Q: The local newspaper writes that Smith "believed the new printing technique would [harm/help] the environment." Do you agree or disagree with this statement?

Responses were collected on a 7-point Likert scale ranging from 1 (completely disagree) to 7 (completely agree).

As discussed above, a mere change in pronouns, by itself, might be insufficient to trigger perspective taking. Participants in the actor conditions thus underwent an additional point-of-view augmentation procedure (POVAP) before the main task, which read:

Imagine you are the chairman of a small company which produces and sells film posters. You have 50 employees, including Frank, your personal assistant who advises you on the business strategy. Your primary goal is to make as much profit as possible for the firm. One day your assistant Frank comes along and says to you: "We have a lot of extra demand for posters, but we cannot print enough posters. If we hire an extra 10 employees we might increase profit." You do your calculations and think to yourself:

“There is a high chance that by hiring the ten extra people, I will increase profit by 20%.” What do you do?

(1) I hire an extra 10 people.

(2) I leave everything as it is.

[New Screen] A few months later your assistant Frank comes to your office and says: “We would have a huge market in Australia. There is a high demand for posters, but few companies produce them. If we want to exploit that opportunity, we should buy a printing factory there, so we would not have to ship the posters all the way across the ocean.” You think to yourself: “With relatively little investment costs, I could double sales and increase profit considerably. The risk that this strategy will fail is only about 10%.” What do you do?

(1) I buy the factory in Australia and start selling posters there.

(2) I leave everything as it is.

Having completed the POVAP, participants faced the target scenario:

One day your assistant Frank comes into your office and says:

“There is a new printing technique which we could use. It is much cheaper and will increase profits, [but/and] it will also [harm/help] the environment.”

You think to yourself: “I don’t care at all about [harming/helping] the environment. I just want to make as much profit as I can, and my calculations show that the new technique will increase profits.”

You decide to implement the new printing technique. As predicted, profits increase and the environment is [harmed/helped].

Q: The local newspaper writes that you “believed the new printing technique would [harm/help] the environment.” Do you agree or disagree with this statement?

Responses were collected on a 7-point Likert scale ranging from 1 (completely disagree) to 7 (completely agree).

Participants who were assigned the observer conditions simply responded to the target scenario (stated above). The rationale for this was that, if undergoing a POVAP analogue, they might also empathize with the chairman, and hence judge from his perspective.⁶ In order to check whether the difference in perspective was manifest, all participants were asked to what extent they had imagined themselves in the position of

⁶ What constitutes responsible priming is a matter of extensive debate. Suffice it to say for now that in the experiments below the topic is addressed in detail. As will be shown, providing participants in the observer conditions with a third-person POVAP analogue does not alter the results.

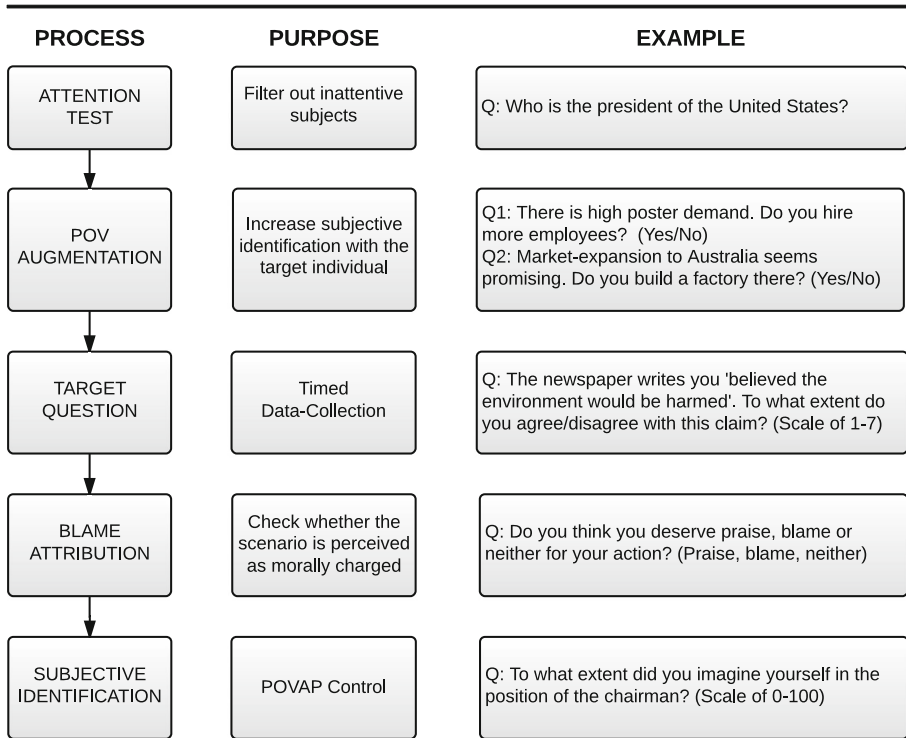


Fig. 1 Survey flow for belief ascriptions (actor perspective, harm condition) in the CHAIRMAN scenario

the chairman on a scale from 0 (not at all) to 100 (completely). Participants were also asked whether they thought the chairman deserved praise, blame or neither. Figure 1 illustrates the survey flow for the actor/harm condition:

3.3 Results & Discussion

A 2(outcome) \times 2(perspective) ANOVA revealed a significant main effect of outcome on belief ascription, $F(1,250)=49.30$, $p<.001$, $\eta_p^2=.165$, no significant main effect of perspective $F(1,250)=.59$, $p=.443$, $\eta_p^2=.002$ and, importantly, a significant interaction $F(1,250)=6.33$, $p=.013$, $\eta_p^2=.025$, which is evidence for a perspective effect.

Consistent with previous findings,⁷ there was a strong ESEE across outcomes in the observer perspective. Belief ratings in the harm condition ($M=6.0$; $SD=2.1$) were significantly higher than in the help condition ($M=3.9$, $SD=1.4$), $t(135)=-6.78$, $p<.001$, the effect size in terms of Cohen's d is large, $d=1.18$ (Cohen 2013). In the actor perspective, belief ratings in the harm condition ($M=5.6$, $SD=1.4$) also differed significantly from those in the help condition ($M=4.6$; $SD=1.8$), $t(115)=-3.22$, $p=.002$. Strikingly, the effect size has been reduced by roughly half from large to medium, $d=.62$. Hence, shifting from an observer to an actor perspective considerably moderates the ESEE (cf. Fig. 2).

⁷ Beebe (2013), cf. also Sverdlík (2004).




Fig. 2 Average belief attribution across outcomes (help v. harm) and perspectives (actor v. observer) in the CHAIRMAN scenario; error bars designate standard error of the mean

A 2(outcome) × 2(perspective) ANOVA revealed a significant main effect of perspective on identification with the chairman, $F(1,250)=67.46, p<.001, \eta_p^2=.213$, no significant main effect of outcome $F(1,250)=1.43, p=.233, \eta_p^2=.006$, and no significant interaction $F(1,250)=.45, p=.504, \eta_p^2=.002$. Aggregating across perspectives, 84% of participants ascribed blame in the harm condition, and 4% in the help condition. The difference in blame ascription was significant across outcomes ($N=254, \chi^2(2)=168.40, p<.001$, Cramer’s $V=.814$, a large effect), though not across perspectives ($N=254, \chi^2(2)=3.65, p=.161$).

In a nutshell, the POVAP seems to work as expected, since identification ratings in the actor perspective significantly exceed those in the observer perspective. Consistent with previous findings, an ESEE was detected across outcome types in the observer perspective. Perspective proved significant: While in the actor perspective, the ESEE remains significant, its size is only about half as pronounced as in the observer perspective. Assuming the point of view of the protagonist in side-effect scenarios thus moderates the Knobe effect for belief ascriptions.

4 Experiment 2: Knowledge in the Chairman Scenario

The second experiment explored the impact of perspective on knowledge ascriptions.

4.1 Participants

Three hundred fifty eight new participants were recruited on Amazon Mechanical Turk to complete a paid online survey. All participants were randomly assigned one of the four conditions of the CHAIRMAN scenario described above. Filtered according to the criteria of Experiment 1, 213 datasets were retained. 113 were from female participants, and the average age of the participants was 40.2 years (SD: 13.0 years).

4.2 Materials and Procedure

The materials and procedure were exactly the same as in Experiment 1, except that the question focused on *knowledge* rather than *belief* ascription. Participants were asked to what extent they agreed that Smith/they themselves knew the new printing technique would harm/help the environment. As before, they also had to report to what extent they identified with the chairman, and whether they thought Smith/they themselves deserved blame, praise or neither for their action.

4.3 Results and Discussion

A 2(outcome) \times 2(perspective) ANOVA revealed a significant main effect of outcome on belief ascription, $F(1,209)=12.18$, $p=.001$, $\eta_p^2=.055$, no significant main effect of perspective $F(1,209)=.07$, $p=.790$, $\eta_p^2=.000$, and no significant interaction $F(1,209)=.36$, $p=.550$, $\eta_p^2=.002$, that is, no perspective effect.

The results replicate previous findings of an ESEE for the observer condition: Knowledge ratings for harm ($M=6.4$; $SD=1.1$) differed significantly from help ($M=5.7$, $SD=1.6$), $t(125)=-3.12$; $p=.002$, $d=0.51$. In the actor condition, the difference between knowledge ratings for harm ($M=6.3$; $SD=0.9$) and those for help ($M=5.7$, $SD=1.5$), was also borderline significant, $t(84)=-1.97$, $p=.052$, $d=0.49$, cf. Figure 3. The effect sizes for the observer and actor conditions were near-identical.

A 2(outcome) \times 2(perspective) ANOVA revealed a significant main effect of perspective on identification with the chairman, $F(1,209)=30.61$, $p<.001$, $\eta_p^2=.128$, no significant main effect of outcome $F(1,209)=.64$, $p=.426$, $\eta_p^2=.003$, and no significant interaction $F(1,209)=.44$, $p=.508$, $\eta_p^2=.002$. Aggregating across actor and observer perspective, 85% of participants ascribed blame in the harm condition, and 1% in the help condition. The difference in blame ascription was significant across outcomes ($N=213$, $\chi^2(2)=156.35$, $p<.001$, Cramer's $V=.857$, a large effect), though not across perspectives ($N=213$, $\chi^2(2)=1.68$, $p=.431$).

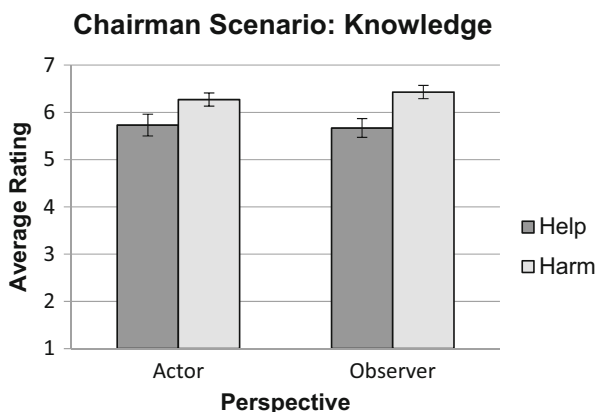


Fig. 3 Average knowledge attribution across outcomes (help v. harm) and perspectives (actor v. observer) in the CHAIRMAN scenario; error bars designate standard error of the mean

Consistent with findings by Myers-Schulz and Schwitzgebel (2013) and Murray et al. (2013), the knowledge ascriptions from Experiment 2 exceed belief ascriptions from Experiment 1. To investigate whether the difference is significant, a three-way ANOVA was conducted with the data of both experiments. There was indeed a significant main effect for state type (belief v. knowledge) on epistemic state ascription ($F(1,459)=47.77$, $p<.001$, $\eta_p^2=.094$). Outcome also proved significant ($F(1,459)=55.21$, $p<.001$, $\eta_p^2=.107$), perspective did not ($F(1,459)=.16$, $p=.687$, $\eta_p^2=.000$). The state type*outcome interaction was significant ($F(1,459)=22.00$, $p=.003$, $\eta_p^2=.019$), as was the perspective*outcome interaction ($F(1,459)=5.06$, $p=.025$, $\eta_p^2=.011$). The state type*perspective interaction was not significant ($F(1,459)=.55$, $p=.459$, $\eta_p^2=.001$), and neither was the three-way interaction ($F(1,459)=2.23$, $p=.136$, $\eta_p^2=.005$). Planned comparisons revealed knowledge to exceed belief ascriptions in all four conditions of the CHAIRMAN scenario (observer/help, $t(135)=5.57$, $p<.001$; observer/harm, $t(125)=2.07$, $p=.040$; actor/help, $t(109)=3.47$, $p=.001$ and observer/harm, $t(90)=2.67$, $p=.009$).

The findings for knowledge ascriptions stand in contrast to those for belief ascriptions. Despite the fact that identification ratings differed strongly across perspectives yet again (suggesting that the POVAP is effective), no significant perspective effect could be detected for knowledge attributions. Previous results of a significant ESEE for the observer perspective were replicated and extended to the actor perspective; the effect sizes for both conditions were very similar.

5 Experiment 3: Knowledge and Belief in the Sales Scenario

Whereas there is a significant impact of perspective on *belief* ascriptions (Experiment 1), there does not seem to be any such impact on *knowledge* ascriptions (no significant interaction of perspective and outcome in Experiment 2). To explore whether these findings are robust across scenarios, another vignette was employed in Experiment 3.

5.1 Participants

For each of the two epistemic states – belief and knowledge – 131 participants were recruited via Amazon Mechanical Turk to complete a paid Qualtrics online survey. Datasets of non-native speakers, subjects who responded in under 20 seconds, and participants failing an attention test or changing their answer ten times or more were excluded. For belief, 115 participants remained, 47 of whom were female. The average age was 39.6 years ($SD=18.2$ years). For knowledge 99 participants remained, 47 of whom were female. The average age was 37.4 years ($SD=11.8$).

5.2 Materials and Procedure

Again following Beebe and Buckwalter (2010) and Beebe and Jensen (2012), an adapted version of Knobe and Mendlow's (2004) SALES scenario was used. The vignette differs from the CHAIRMAN scenario in so far as the agent produces a negative side-effect, which is not clearly morally bad or in conflict with a salient norm. In previous studies, the agent was thus not considered to be blameworthy in the negative

side-effect condition (Phelan and Sarkissian 2008). Participants were randomly assigned to one of the eight conditions (2 perspectives x 2 outcomes x 2 epistemic state types). The first-person negative outcome condition focusing on belief read:

Imagine you are the chairman of a small company which produces and sells film posters. You have 50 employees, including Frank, an assistant who advises you on the business strategy.

One day your assistant Frank comes into your office and says to you: “We are thinking of implementing a new corporate restructuring plan. It will simplify our corporate structure but it will also decrease sales in New Jersey.”

You respond to Frank: “I don’t care about what happens in the next quarter. We need to simplify our corporate structure. Let’s implement the new plan.”

You implement the new plan. Sales in New Jersey decrease.

Q: The local newspaper writes that you “believed that the plan would decrease sales in New Jersey”. Do you agree or disagree with this statement?

For the third-person condition, ‘you’ was replaced with ‘Mr Smith’, for the positive outcome condition ‘decrease’ was replaced with ‘increase’, and for the knowledge condition ‘believed’ was replaced with ‘knew’. Responses were collected on a 7-point Likert scale ranging from 1 (completely disagree) to 7 (completely agree). Participants were again asked to what extent they identified with the chairman of the company on a scale from 0 (not at all) to 100 (completely), and whether they thought the chairman deserved praise, blame or neither. Actor perspective participants underwent a POVAP with two managerial decisions, participants assigned to the observer conditions did not.

5.3 Results

5.3.1 Three-Way ANOVA

A 2(outcome) × 2(perspective) × 2(epistemic state type) ANOVA revealed a significant main effect of outcome on epistemic state ascription, $F(1,206)=42.86, p<.001, \eta_p^2=.172$; a significant main effect of epistemic state type, $F(1,206)=15.18, p=.013, \eta_p^2=.029$ and no significant main effect of perspective $F(1,206)=.04, p=.842, \eta_p^2=.000$. The outcome*perspective interaction was significant, $F(1,206)=5.29, p=.022, \eta_p^2=.025$, the outcome*state type interaction was not, $F(1,206)=.03; p=.854, \eta_p^2=.000$ and neither was the perspective*state type interaction, $F(1,206)=1.302, p=.255, \eta_p^2=.006$. The three-way interaction was significant, $F(1,206)=4.34, p=.039, \eta_p^2=.021$.

5.3.2 Results for Belief Ascriptions

A 2(outcome) × 2(perspective) ANOVA revealed a significant main effect of outcome on belief ascription, $F(1,111)=26.07, p<.001, \eta_p^2=.190$, no significant main effect of perspective $F(1,111)=.57, p=0.452, \eta_p^2=.005$, and a significant interaction

$F(1,111)=12.37, p=.001, \eta_p^2=.100$, i.e. a significant perspective effect. Consistent with previous studies, there was a strong ESEE across outcomes for the observer perspective. Belief ratings in the decrease condition ($M=5.8; SD=1.1$) were significantly higher than in the increase condition ($M=3.4, SD=1.7$), $t(55)=-6.16, p<.001; d=1.68$, a very large effect. As regards the actor perspective, belief ratings in the decrease condition ($M=5.0; SD=1.5$) and increase condition ($M=4.6, SD=1.5$) did not differ significantly, $t(56)=-1.114; p=.270, d=.27$. The first-person perspective crowds out the third-person perspective ESEE entirely, cf. Figure 4.

As regards identification with the protagonist, a $2(\text{outcome}) \times 2(\text{perspective})$ ANOVA revealed a predicted main effect of perspective, $F(1,111)=32.29, p<.001, \eta_p^2=.225$, no significant main effect of outcome $F(1,111)=.29, p=.590, \eta_p^2=.003$, and no significant interaction $F(1,111)=.05, p=.829, \eta_p^2=.000$.

Aggregating across perspectives, 38% of participants ascribed blame in the negative side-effect conditions, and 4% in the positive side-effect conditions. The difference in blame ascriptions was significant across outcomes ($N=115, \chi^2(2)=22.505, p<.001$, Cramer's $V=.442$, a medium effect) but not across perspectives ($N=115, \chi^2(2)=1.877, p=.391$).

5.3.3 Results for Knowledge Ascriptions

A $2(\text{outcome}) \times 2(\text{perspective})$ ANOVA revealed a significant main effect of outcome on knowledge ascription, $F(1,95)=17.80, p<.001, \eta_p^2=.158$, no significant main effect of perspective $F(1,95)=.71, p=.403, \eta_p^2=.007$, and no significant interaction $F(1,95)=.02, p=.891, \eta_p^2=.000$.

Replicating previous findings, there was a pronounced ESEE for the observer perspective: Knowledge ratings in the decrease condition ($M=6.2; SD=1.3$) were significantly higher than in the increase condition ($M=4.6, SD=1.7$), $t(32)=-2.89, p=.007; d=1.06$, a large effect. In the actor perspective, knowledge ratings for the decrease condition ($M=5.8, SD=1.4$) also differed significantly from the increase condition ($M=4.4; SD=2.0$), $t(63)=-3.37; p=.001; d=.60$, a medium effect, cf. Figure 5.

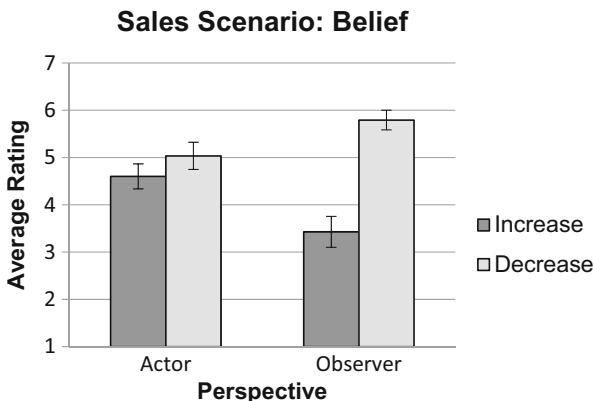


Fig. 4 Average belief attribution across outcomes (increase, decrease) and perspectives (actor, observer) in the SALES scenario; error bars designate standard error of the mean

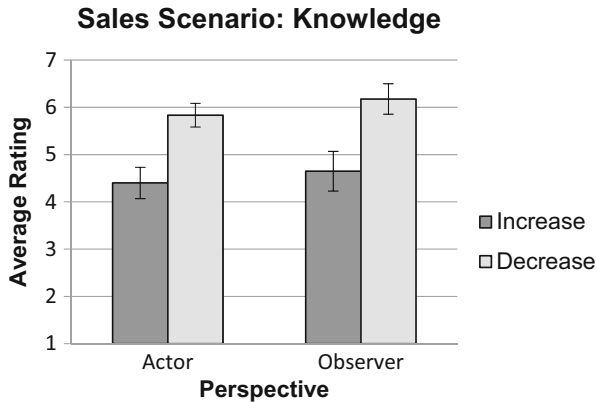


Fig. 5 Average knowledge attribution across outcomes (increase, decrease) and perspectives (actor, observer) in the SALES scenario; error bars designate standard error of the mean

A 2(outcome) \times 2(perspective) ANOVA for reported identification with the chairman revealed a significant main effect of perspective, $F(1,95)=40.18$, $p<.001$, $\eta_p^2=.297$, suggesting that the POVAP was effective, no significant main effect of outcome $F(1,95)=.00$, $p=.981$, $\eta_p^2=.000$, and no significant interaction $F(1,95)=1.50$, $p=.223$, $\eta_p^2=.016$.

Aggregating across perspectives, 19% of participants ascribed blame in the negative side-effect condition, and 0% in the positive side-effect condition. The difference in blame ascriptions was significant across outcomes ($N=99$, $\chi^2(2)=24.00$, $p<.001$, Cramer's $V=.492$, a medium effect) but not across perspectives ($N=99$, $\chi^2(2)=1.74$, $p=.410$).

For two conditions, the level of mean knowledge ascription significantly exceeded the level of mean belief ascription (actor/decrease: $t(56)=-2.10$, $p=.040$; observer/increase: $t(43)=-2.29$, $p=.027$). For the remaining conditions, no significant difference could be detected: actor/increase: $t(63)=.46$, $p=.646$; observer/decrease: $t(44)=-1.05$, $p=.301$).

5.4 Discussion

The difference in identification with the vignette's protagonist across perspectives once again testifies to the efficacy of the POVAP: Identification in the actor conditions significantly exceeds identification in the observer conditions. Replicating previous findings, there is a significant ESEE across outcomes for both knowledge and belief in the observer condition. Just as in Experiments 1 and 2, perspective proved significant for belief ascriptions, but not for knowledge ascriptions.

Though perspective had an effect on belief attribution in both the CHAIRMAN and the SALES scenario, the findings differ in one important respect: While in the CHAIRMAN scenario the ESEE is merely moderated in the actor perspective, in the SALES scenario it is cancelled out entirely. Since the scenarios differ principally with respect to the salient moral norms and – presumably relatedly – the perceived blameworthiness of the protagonist, one might speculate about a relationship between norm infractions and/or blame, the ESEE and perspective. When norm infractions and blame do not play an important role, such as in SALES, the first-person perspective crowds out the ESEE

entirely. When norm infractions and blame do play a role, such as in *CHAIRMAN*, they might reduce the effect of perspective taking, perhaps because it is more difficult to put oneself in the shoes of a dislikable character willing to violate what we consider salient norms.

5.5 Methodological Concerns

The experiments presented so far might raise two concerns. The first one regards perspective and knowledge ascriptions. The absence of evidence for a significant difference across actor and observer perspectives in knowledge ascriptions does not constitute evidence in favour of the absence of a significant difference in perspective. The sample size for the knowledge condition of Experiment 3 might be deemed too low to rule out Type II errors. Experiment 4 attempts to replicate the findings for knowledge with a larger sample size determined on the basis of an *a priori* power calculation.

A second worry regards the POVAP. As discussed, previous experiments have shown that a mere change in pronouns is insufficient to trigger perspective taking. However, the fact that participants in the actor condition undergo the point of view augmentation procedure, while those in the observer condition do not receive a third-person POVAP analogue (i.e. a POVAP invoking only third-person pronouns) might go beyond acceptable priming. The difference in results between the POVAP-preceded actor tasks on the one hand, and the observer tasks on the other, might arise not due to perspective-taking, but due to the fact that the former group spends more time and thought on the materials and thus arrives at a more considered verdict. In a nutshell, the concern would be that what appears as a perspective effect might instead be the impact of reflective judgment.

I do not consider this line of reasoning very convincing, for the following reasons. First, one would presumably have to assume that the consistently higher identification ratings with the protagonist in the actor conditions are driven principally by a cooperation bias. While not impossible, this does not seem particularly plausible. Second, and more importantly, if the difference across actor/observer conditions for belief ascriptions were indeed driven exclusively by more careful reflection, we would expect the results of knowledge ascription to pattern with those for belief ascription. However, for knowledge the hypothesized more extensive deliberation in the actor condition does not leave a mark on the ascription of the mental state. Third, the impact of reflective judgment, currently a topic of much debate in meta-philosophy, is likely minimal. Critics of experimental philosophy argue that one reason why survey-based folk-intuitions are of little import to philosophical theorizing is that the collected intuitions are crude shots from the hip, rather than the product of careful and nuanced deliberation characteristic of philosophy (Kauppinen 2007; Ludwig 2007). However, a flurry of recent work suggests that extensive deliberation simply does not influence folk judgment.⁸

All this being said, the remaining studies attempt to resolve the matter empirically by including a POVAP analogue in the observer conditions. Following the replication of the

⁸ Colaço, Kneer, Alexander & Machery, On second thought: A refutation of the reflection defense (ms.) employ a number of manipulations used in social psychology and experimental economics to increase deliberation. In a series of five experiments testing classic philosophical thought-experiments, not a single manipulation had a significant impact on judgment. Similarly, studies by Schwitzgebel and Cushman (2012, 2015) find that the judgments of laymen and trained philosophers, with their penchant for careful and extensive deliberation, manifest very little difference. For further discussion of reflective judgment, cf. Alexander and Weinberg (2007), Weinberg et al. (2012).

knowledge condition for the SALES scenario with a larger sample size, a third scenario familiar from the ESEE literature was run for belief and knowledge ascriptions.

6 Replicating the Results of the Knowledge Experiments

The main goal of Experiment 4 was to test whether the absence of evidence of a perspective effect for knowledge ascriptions would prove robust for a larger sample size. To address the worry regarding a potential difference in reflective deliberation across perspectives, participants in the observer condition were also presented with a POVAP analogue discussed in detail below.

6.1 Participants

In the belief condition of Experiment 3, we found a perspective effect of $\eta_p^2=.100$ (equivalent to Cohen's $f=.33$). According to power calculations with *G*Power* (Faul et al. 2009) using standard levels of $\alpha=.05$ and $(1-\beta)=.80$, a sample size of 103 participants – near-exactly what was used in Experiment 3 – should be sufficient to detect a similar medium-sized perspective effect for knowledge ascriptions. In order to detect a small perspective effect of $f=.2$ on the same assumptions, we would require a sample size of 277 subjects. Differently put, on standard assumptions, a sample size of 277 subjects should suffice to produce a perspective effect of a considerably smaller magnitude than what we found for belief ascriptions. So as to ensure that the required number of participants could be met after excluding non-native speakers, inattentive subjects and the like, 404 participants were recruited on Amazon Mechanical Turk. Datasets of non-native speakers, subjects who responded in under 20 seconds, and participants failing an attention test or changing their answer ten times or more were excluded. 251 datasets were retained, of which 130 were from female participants. The average age was 37,1 years ($SD=11,2$).

6.2 Materials and Procedure

The experiment took a 2(outcome: decrease v. increase) \times 2(perspective: actor v. observer) design and was identical with the knowledge condition of Experiment 3, except in one respect: All participants in the observer condition underwent a third-person analogue to the POVAP completed by those in the actor condition. That is, observer participants also faced the vignettes regarding managerial decisions focusing on the potential hiring of further employees and business expansion overseas. This ensured that all participants had read the same amount of text, spent a similar period of time with the materials, and processed similar contents.

6.3 Results and Discussion

A 2(outcome) \times 2(perspective) ANOVA revealed a significant main effect of outcome on knowledge ascriptions, $F(1,247)=39.10$, $p<.001$, $\eta_p^2=.137$, no significant main effect of perspective, $F(1,247)=2.13$, $p=.146$, $\eta_p^2=.009$, and – the main focus of the replication – no significant outcome*perspective interaction, $F(1,247)=1.00$, $p=.318$, $\eta_p^2=.004$. The results of Experiment 2 and 3, according to which no perspective effect (i.e. no significant

interaction) can be detected for knowledge ascriptions, was thus replicated. Consistent with the literature, there was a significant ESEE for the observer perspective (decrease: $M=6.3$, $SD=1.1$; increase: $M=5.0$, $SD=1.7$, $t(128)=5.20$, $p<.001$). Consistent with the findings reported above, there was also a significant ESEE for the actor perspective (decrease: $M=5.9$, $SD=1.1$, increase: $M=4.9$, $SD=1.6$, $t(119)=3.67$, $p<.001$).

A 2(outcome) \times 2(perspective) ANOVA for self-reported identification revealed a significant main effect of perspective, $F(1,246)=59.20$, $p<.001$, $\eta_p^2=.194$; no significant main effect of condition, $F(1,246)=.49$, $p=.484$, $\eta_p^2=.002$, and no significant interaction, $F(1,246)=2.44$, $p=.120$, $\eta_p^2=.010$. As regards the decrease condition, identification in the actor perspective ($M=81.5$, $SD=22.7$) significantly exceeded identification in the observer perspective ($M=61.7$, $SD=28.5$), $t(124)=4.28$, $p<.001$. As regards the increase condition, identification in the actor perspective ($M=88.8$, $SD=14.7$) also significantly exceeded identification in the observer perspective ($M=58.9$, $SD=31.8$), $t(122)=6.64$, $p<.001$. If self-reported identification is a meaningful indicator for perspective-taking, the POVAP proves once again effective, and the POVAP analogue does not distort the control condition.

In sum: A pronounced difference in perspective taking can be inferred from the reported identification with the scenario's protagonist across actor and observer conditions. This suggests that an appropriately implemented POVAP analogue does not distort the observer condition. Replicating the above reported findings, knowledge ascriptions manifest a significant ESEE, and we once again failed to detect a significant perspective effect (the perspective*condition interaction), despite using a large sample size.

7 Experiment 5: Knowledge and Belief in the Movie Scenario

In order to explore whether the findings so far reported are robust, a final experiment employed another vignette familiar from the ESEE literature, the MOVIE scenario (Beebe and Jensen 2012, adapted from Knobe 2004a). The experiment design paid close attention to the two caveats discussed at the end of section 5. First, to address concerns regarding reflective judgment, participants in the observer condition received an analogue of the POVAP. Second, to address sample size worries for the knowledge condition, an a priori power calculation was conducted on the basis of the effect size of the perspective effect for belief ascriptions.

7.1 Participants for the Belief Condition

For the belief condition, 205 participants were recruited on Amazon Mechanical Turk to complete a Qualtrics online survey. Datasets of non-native speakers, subjects who responded in under 20 seconds, and participants failing an attention test or changing their answer ten times or more were excluded. 144 datasets remained, 87 of which were from female participants. The average age was 41,1 years ($SD = 13,6$).

7.2 Participants for the Knowledge Condition

A pilot ($N=100$) for the belief condition produced a perspective effect of $\eta_p^2=.046$ (\approx Cohen's $f=.22$), a small effect size. Using standard levels of $\alpha=.05$ and $(1-\beta)=.80$ as

well as the pilot's effect size for belief, calculations with *G*Power* produced a recommended sample size of 231. Differently put, a sample size of 231 should suffice to produce a small perspective effect of a magnitude equivalent to what can be expected for belief ascriptions under standard assumptions for α and β . So as to ensure that the required number of participants could be met after excluding non-native speakers, inattentive subjects and the like, 401 participants were recruited on Amazon Mechanical Turk. Applying the same exclusion criteria as for the belief condition, 261 datasets were retained, 126 of which were from female participants. The average age of the participants was 36,4 years ($SD = 12,1$).

7.3 Materials and Procedure

The experiment took a 2(epistemic state: belief v. knowledge) \times 2(outcome: better films v. worse films) \times 2(perspective: actor v. observer) design. In contrast to Experiments 1 and 3, all subjects underwent the POVAP or a third-person analogue. The sole difference between the POVAP and its analogue consisted in first-person pronouns v. third-person pronouns (or 'Mr Smith'/'the CEO'). To take an example, the belief/worse films/actor condition read:

Imagine you are the CEO of a movie production company. Your company makes comedies, action films, dramas and science fiction movies. You have 90 employees, including Mr Jones, who advises you on the company's business strategy. Your goal is to make as much profit as possible for the company. One day, Jones says to you: "The comedies we produce are twice as profitable as our action movies." You think to yourself: "If we used our budget to make more comedies and less action movies, we would make more profit."

Q.: What will you do?

- (1) I produce more comedies and less action movies.
- (2) I leave everything as it is.

[New screen] On another occasion, Jones, your strategy advisor, says to you: "Demand for our films is very high. We might want to think about hiring more people." You think to yourself: "If we hired five new producers, we could complete an extra ten films per year."

Q: What will you do?

- (1) I hire new producers.
- (2) I leave everything as it is.

[New Screen] A few months later, Jones, the strategy advisor, says to you: "We are thinking of implementing a new policy. If we implement the policy, it will increase profits for our company, but it will also make our movies worse from an

artistic standpoint.” You think the proposal over and reply to Jones: “Look, I am not concerned about how good our movies are from an artistic point of view. My task is to produce as much profit as possible for the company. Let’s implement the new policy.” You have the new policy implemented. The policy does indeed make the movies worse from an artistic standpoint.

Q: The local newspaper writes that you “believed the new policy would make the movies worse from an artistic standpoint.” To what extent do you agree or disagree with this claim?

In the observer conditions, the pronouns ‘you’ in the text, and ‘I’ in the questions and responses, were replaced with ‘Mr Smith’, ‘he’ or ‘the CEO’. In the positive outcome conditions, ‘worse’ was replaced with ‘better’; in the knowledge conditions, ‘believed’ was replaced with ‘knew’.

Responses were collected on a 7-point Likert scale ranging from 1 (completely disagree) to 7 (completely agree). Participants were also asked to what extent they identified with the chairman on a 7-point Likert scale ranging from 1 (not at all) to 7 (completely), and whether they thought they deserved blame for their action on a 7-point Likert scale ranging from 1 (no blame at all) to 7 (a lot of blame).

7.4 Results

7.4.1 Three-Way ANOVAs

A 2(outcome) × 2(perspective) × 2(epistemic state type) ANOVA revealed a significant main effect of outcome on epistemic state ascriptions, $F(1,397)=217.19$, $p<.001$, $\eta_p^2=.198$, a significant main effect of state type, $F(1,397)=49.24$, $p<.001$, $\eta_p^2=.110$, and no significant main effect of perspective $F(1,397)=.032$, $p=.857$, $\eta_p^2=.000$. The outcome*state type interaction was significant, $F(1,397)=26.84$; $p=.001$, $\eta_p^2=.030$. The outcome*perspective interaction was also significant, $F(1,397)=6.81$, $p=.009$, $\eta_p^2=.017$, the perspective*epistemic state type interaction was not, $F(1,397)=1.83$, $p=.178$, $\eta_p^2=.005$. The three-way interaction was insignificant, $F(1,397)=3.69$, $p=.055$, $\eta_p^2=.009$.

A second three-way ANOVA revealed a significant main effect of perspective on identification with the protagonist $F(1,397)=84.52$, $p<.001$, $\eta_p^2=.176$. There were no significant main effects of epistemic state type $F(1,397)=.38$, $p=.540$, $\eta_p^2=.001$ or outcome $F(1,397)=.713$, $p=.399$, $\eta_p^2=.002$. None of the interactions were significant (all $ps>.252$).

7.4.2 Results for Belief Ascriptions

A 2(outcome) × 2(perspective) ANOVA revealed a significant main effect of outcome on belief ascriptions, $F(1,140)=54.50$, $p<.001$, $\eta_p^2=.280$, no significant main effect of perspective, $F(1,140)=.519$, $p=.402$, $\eta_p^2=.003$ and a significant outcome*perspective interaction, $F(1,140)=6.25$, $p=.014$, $\eta_p^2=.043$.

In the observer condition, there was a significant ESEE across outcome types (worse: $M=6.0$, $SD=1.4$; better: $M=3.2$, $SD=1.8$; $t(73)=7.54$, $p<.001$, $d=1.72$). In the

actor condition, the ESEE was also significant, though its effect size was less than half of the ESEE for the observer condition (worse: $M=5.5$, $SD=1.3$; better $M=4.1$, $SD=2.1$, $t(67)=3.20$, $p=.002$, $d=0.78$).

In an ANOVA focusing on blame ascriptions, outcome proved significant, $F(1,140)=29.92$, $p<.001$, $\eta_p^2=.176$, as did perspective, $F(1,140)=4.38$, $p=.038$, $\eta_p^2=.030$ and the outcome*perspective interaction, $F(1,140)=6.16$, $p=.014$, $\eta_p^2=.042$.

7.4.3 Results for Knowledge Ascriptions

A $2(\text{outcome}) \times 2(\text{perspective})$ ANOVA revealed a significant main effect of outcome on knowledge ascriptions, $F(1,257)=34.06$; $p<.001$, $\eta_p^2=.117$. Perspective was insignificant, $F(1,257)=1.94$, $p=.165$, $\eta_p^2=.007$, as was the perspective*outcome interaction, $F(1,257)=.39$, $p=.532$, $\eta_p^2=.002$. In the actor condition, there was a significant ESEE across outcome types (worse: $M=6.1$, $SD=1.2$; better: $M=5.2$, $SD = 1.7$; $t(133)=3.49$, $p=.001$, $d=.60$). In the observer condition, the ESEE was also significant (worse: $M=6.4$, $SD=.9$; better $M=5.3$, $SD=1.6$, $t(124)=4.93$, $p<.001$, $d=.86$).

An ANOVA focusing on blame ascriptions revealed a significant main effect of outcome, $F(1,257)=84.77$, $p<.001$, $\eta_p^2=.248$. Perspective was not significant, $F(1,257)=.13$, $p=.724$, $\eta_p^2=.000$, and neither was the outcome*perspective interaction, $F(1,257)=1.88$, $p=.171$, $\eta_p^2=.007$.

For all four conditions, the level of mean knowledge ascriptions significantly exceeds the level of mean belief ascriptions (actor/harm: $t(97)=-2.45$, $p=.016$; actor/help: $t(103)=-2.92$, $p=.004$, observer/harm: $t(107)=-2.17$, $p=.032$, observer/help: $t(90)=-5.89$, $p<.001$).

In short, all results from Experiments 1–4 were replicated. In the observer conditions, there is a significant ESEE for both knowledge and belief ascriptions. For belief ascriptions, the ESEE is reduced to less than half when participants judge from a first-person perspective; for knowledge ascriptions, no significant perspective effect could be detected.

8 General Discussion

8.1 Summary of Findings

The experiments here reported produced four main findings. One is methodological, the other three are substantive. First, and *pace* Feltz et al. (2012), the methodology of using imaginary scenarios is not inherently unsuited to the study of perspective and perspective taking. Second, and replicating the studies of Beebe and Buckwalter (2010), Beebe and Jensen (2012) and Beebe (2013), the ESEE is robust for epistemic state ascriptions in the observer perspective across all three scenarios. Third, there is a perspective effect for belief ascriptions. Since perspective works in the opposite direction of the ESEE, it cancels out the ESEE either partly (CHAIRMAN, MOVIES) or wholly (SALES). Finally, in contrast to belief, no evidence in favour of a perspective effect for knowledge ascriptions was found. I will briefly discuss these findings in turn.

8.2 Methodology: Surveys versus Active Participation

Feltz et al. (2012) ran online surveys similar to the ones presented here, though their dependent variable was intentionality rather than epistemic states. The difference in perspective consisted in a mere variation of pronouns ('you, the chairman' v. 'he, the chairman'). No significant effect due to perspective was found. However, in a laboratory experiment, where participants engaged in an authentic (rather than imaginary) investment game, intentionality ascriptions of those participating in the game and those observing the players came apart. Feltz et al. conclude that 'actor-observer differences are not likely to be found by simply asking participants to imagine that they are actors' and that bringing out perspectival asymmetries requires 'putting participants in the actual decision making environment' (2012, p.682–3).

As demonstrated, this assessment is not quite correct. Though a simple substitution of third- to second-person pronouns might indeed not do the trick, participants can be incited to imagine themselves in the perspective of a scenario's protagonist, for instance by means of the here employed point-of-view augmentation procedure. The reported identification with the protagonist differed significantly across actor/observer perspectives for all experiments. This finding also holds for Experiment 4 and 5, where a third-person POVAP analogue was employed to rule out the possibility that the POVAP merely increases engagement with the target scenario and thus leads to more reflective judgments instead of perspective taking. The identification ratings suggest that perspective augmentation works, and hence that the decrease in asymmetric belief ascriptions across the actor/observer conditions are in fact due to perspective. The empirical exploration of perspective effects is hence not limited to active participation in laboratory experiments or real-life situations.

8.3 The ESEE and Doxastic Heuristics

The replication of the ESEE for both belief and knowledge in the observer condition lends some support to Alfano et al.'s (2012) theory, which sees the ESEE for belief at the root of the Knobe effect for all mental state ascriptions. However, some of our results also cast doubt on the account.

First, if the side-effect effect for mental states were always driven by the side-effect effect for belief, knowledge ascriptions should manifest a similar pattern across perspective as belief ascriptions. However, they do not, which suggests that despite the close conceptual link between the concepts of belief and knowledge, their *ascriptions* are sensitive to different factors.

Second, recall that on the view of Alfano and colleagues, other-ascriptions of mental states are guided by a rule of thumb – the norm-violation belief-ascription (NBA) heuristic – according to which we should attribute beliefs to others regarding negative side-effects, but not positive ones. NBA is hypothesized to be an extension of a rule to which we adhere in self-ascriptions of belief (NBF): If my action gives rise to a negative side-effect, I have good reason to reflect on, and hence form a belief about, said side-effect. This neat explanatory chain comes under pressure from the data on belief ascriptions. In the CHAIRMAN and MOVIE scenarios, asymmetric belief attribution is significantly reduced when participants evaluate themselves rather than another person. In the SALES scenario, there is no asymmetric pattern in belief self-attribution

(no ESEE) despite the differences in norm-conformity across conditions. It is thus not evident whether the folk follow the NBF heuristic ('If my own φ -ing would make it the case that p and p violates a norm salient to me, believe that φ -ing would make it the case that p .')⁹ But if they do not, and if NBF is to justify NBA, the whole theory comes under pressure.

8.4 The Perspective Effect

In the psychological literature concerning the effect of perspective taking on phenomena such as the appreciation of others, forgiveness, moral judgment etc. (cf. section 1.1), perspective taking itself is standardly treated as explanatorily irreducible. Whether this is an adequate approach is an interesting and complex question. One might, for instance, explore whether perspective effects of the sort here reported can be subsumed under the self-serving bias, according to which individuals' assessment of their own abilities and actions are consistently biased in their favour. In the particular case at hand, this makes good intuitive sense: A more pronounced tendency to self-attribute belief about positive side-effects and refusal to do so for negative side-effects portrays the agent favourably in comparison to other-assessment. A detailed treatment of the explanatory role and potential reducibility of perspective and perspective taking to more basic phenomena would take us too far afield. However, it might be helpful to explore whether the effect detected for belief ascriptions does indeed capture a genuine impact of perspective taking (explanatorily basic or not), or whether it merely appears to do so.

Taking a cue from the standard explanations of the actor/observer asymmetry,¹⁰ one might, for instance, question whether (i) in the actor position, the subject had *better epistemic access* to her reasons (suggesting a possibly unwanted asymmetry across scenarios). Or else, one might wonder whether (ii) participants in the actor conditions merely wanted to *present* themselves in a favourable light, suggesting that what was in fact reported was distinct from what was apparently measured.

Though plausible for actor/observer experiments properly understood,¹¹ the explicit and forced identical stipulation of the protagonist's reasons across perspectives makes asymmetric access to reasons a rather unconvincing explanation of the present findings. Reputation management, however, seems promising at first glance: Participants in the actor position might shy away from assuming responsibility for the negative consequences through lower self-ascription of belief, and present themselves in a favourable light by higher self-ascription of belief in the positive side-effect scenarios. However, two reasons speak against this explanation: Firstly, if reputation management were an important concern, we would have likely seen a similar pattern for knowledge ascriptions as for belief ascriptions. Secondly, if participants had intended to make a positive impression, this should have manifested itself clearly at the level of blame and praise

⁹ Note that blame ascriptions differed significantly across conditions for all three scenarios, which suggests that the normative differences were indeed salient to the participants. There was no significant difference for blame ascriptions across perspectives, which blocks a defence of Alfano et al.'s account according to which the actor perspective reduces the *salience* of norm-violation (cf. for instance, Robinson et al. 2015).

¹⁰ Cf. Malle et al. (2007).

¹¹ The term *actor/observer asymmetry* should be used with care, as it was originally intended to refer to the difference in explanation *types* regarding the behaviour of oneself v. others (cf. Jones and Nisbett 1971), rather than perspective effects broadly conceived.

ascriptions. However, blame ascriptions in the belief conditions did not differ significantly across perspectives for either CHAIRMAN or SALES and only for MOVIES could a small effect of perspective on blame be detected. Overall, it thus appears unlikely that the difference in belief ascription across perspectives is merely due to a difference in the way the situation is *reported*. The detected perspective effect seems to be driven by the way the situation is *perceived* and *judged*.

8.5 Belief versus Knowledge Ascriptions

At the outset of our inquiry into the effects of perspective on belief and knowledge ascriptions, we discussed four possible hypotheses: Perspective might affect both knowledge and belief attributions, neither, or else only one of the two types of states. It turned out that belief attributions do manifest a perspective effect, though knowledge attributions do not. For traditional epistemologists, this might come as a surprise. Knowledge is taken to entail belief, and so one might expect the two types of epistemic state to respond similarly to perspective taking. One way to accommodate the findings consists in breaking with the entailment thesis (and hence with traditional epistemology): If knowledge were conceived as conceptually independent of belief, as Williamson argues and recent data by Myers-Schulz & Schwitzgebel seems to suggest,¹² there is no reason to presume that knowledge and belief must be affected in the same way by perspective taking.

I find the entailment thesis plausible (though will refrain from arguing the case here).¹³ I'd thus like to propose an explanation invoking epistemic authority, which does not require knowledge to be conceived as a *sui generis* mental state, and which is consistent with the assumptions of orthodox epistemology. As briefly discussed in the introduction, there is an asymmetry in epistemic authority across self-ascriptions of belief on the one hand, and knowledge on the other. Whether or not I do indeed believe that *p* – complications pertaining to the subconscious or ‘being in two minds’ aside – can only authoritatively be assessed by me. Whether or not I *know* that *p*, however, depends inter alia on whether *p* is true – a matter on which I (except as regards certain types of *de se* knowledge) enjoy no more authority than anyone else. While I am hence in a very privileged position when it comes to assessing my own beliefs, I am not the sole and ultimate authority concerning which of my beliefs do, in fact, constitute knowledge.

Linguistically, the difference is very salient. If I say ‘I know it’s snowing’ you can respond with ‘No, it isn’t’ or ‘No, you don’t’. The first response challenges the proposition embedded under the attitude verb; the second one challenges my being in the very epistemic state I take myself to be in. To do the latter, it suffices to have good grounds to challenge the proposition embedded. If, however, I say ‘I believe it’s snowing’, the response ‘No you don’t’ is standardly infelicitous, no matter whether it is actually snowing. My self-report of *having* the stated belief attitude cannot easily be contradicted by anyone, only the embedded proposition can. Among the four subtypes of epistemic state-attribution, i.e. self- and other-attribution of belief, and self- and

¹² In fact, our results are broadly consistent with such an unorthodox view of epistemology: For half of the conditions of the SALES experiment, and all of the four conditions of the CHAIRMAN and MOVIES experiments, mean knowledge ascriptions significantly exceed mean belief ascriptions.

¹³ For an interesting discussion of the results of Myers-Schulz and Schwitzgebel, see Rose and Schaffer (2013) and Buckwalter et al. (2015).

other-attribution of knowledge, self-ascription of belief is thus distinctive as it is characterized by special epistemic privileges. In perspective *taking*, we seem to avail ourselves of similar privileges. Note that they are not strictly warranted: Our epistemic situation in judging a target subject barely differs across an observer and an assumed first-person perspective. But when inhabiting another's point of view, we take ourselves to enjoy the same privileged position, and hence the same liberties as in belief-ascription from our own point of view. This suggests that belief ascriptions are responsive not only to the available *evidence* in favour of, or against, attributing a particular belief state to others (such evidence, in our scenarios, was held fixed), but are also sensitive to the *mode* (first- v. third-person) from which the judging subject undertakes the ascription. All this is consistent with the entailment thesis in particular, and traditional epistemology more generally. Given that our epistemic authority is, and is perceived as, more limited in knowledge self-ascriptions than in belief self-ascriptions, the sense of epistemic privilege does not spread from genuine to imagined self-ascriptions of knowledge as it does for belief.

8.6 Future Research

Experimental philosophy near-uniformly neglects the first-person perspective, and the related, yet distinct, phenomenon of perspective-taking has received equally little attention. This is a considerable shortcoming of the discipline, and it is likely that many a theory drawn up on the basis of observer-perspective evidence falters with respect to actor-perspective data. If the goal of the discipline is to elucidate central folk concepts such as knowledge, belief, intention, causation and the like in general, it stands to reason to move beyond the observer-fixated experimental paradigm.

The studies here presented are modest in ambition, and focus on the *imagined* first-person perspective, rather than the first-person perspective *proper*. Future research should, where appropriate and possible, opt for a tri-partite approach that investigates the phenomenon under consideration for the genuine first-person perspective, the *imagined* first-person point of view and the third-person perspective. A revised experimental paradigm of this sort would not only produce data that is informative across perspectives, but allow for a systematic exploration of the similarities and differences between judgments from one's own perspective and judgments from an assumed target perspective.

While the point-of-view augmentation procedure here employed has rendered us good service for the inquiry into perspective taking, there is ample room for improvement. For instance, it is difficult to assess to what extent perspective taking in the observer condition is *also* augmented by the POVAP analogue employed in Experiments 4 and 5, and to what extent the effect sizes reported are accurate. Ideally, a broad variety of measures that foster perspective taking should be tested, so as to determine what manipulations are best suited to the task.

The astonishing asymmetric impact of perspective on belief and knowledge ascriptions was explained with regards to a difference in epistemic authority across different types of epistemic states. While tentative at best for now, it would be interesting to further investigate the hypothesis empirically and to explore its theoretical implications in depth.

Acknowledgements Thanks to Mark Alfano, James Beebe, Wesley Buckwalter, Quassim Cassam, Florian Cova, Simon Cullen, Stefanie Kneer, Edouard Machery, Barry Maguire, Blake Myers-Schulz, Obioma Ofoego, Brian Robinson, David Rose, John Turri and three anonymous reviewers for their insightful comments. I am particularly grateful to Joshua Knobe and Brent Strickland for very helpful feedback.

References

- Alexander, J., and J.M. Weinberg. 2007. Analytic epistemology and experimental philosophy. *Philosophy Compass* 2 (1): 56–80.
- Alfano, M., J.R. Beebe, and B. Robinson. 2012. The centrality of belief and reflection in Knobe-effect cases: A unified account of the data. *The Monist* 95 (2): 264–289.
- Batson, C.D., J.H. Eklund, V.L. Chernok, J.L. Hoyt, and B.G. Ortiz. 2007. An additional antecedent of empathic concern: Valuing the welfare of the person in need. *Journal of Personality and Social Psychology* 93 (1): 65–74.
- Beebe, J.R. 2013. A Knobe effect for belief ascriptions. *Review of Philosophy and Psychology* 4 (2): 235–258.
- Beebe, J.R. 2016. Evaluative effects on knowledge attributions. In *A Companion to Experimental Philosophy*, ed. J. Sytsma and W. Buckwalter. Hoboken: John Wiley & Sons.
- Beebe, J.R., and W. Buckwalter. 2010. The epistemic side-effect effect. *Mind & Language* 25 (4): 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 (5): 689–715.
- Beebe, J.R., and J. Shea. 2013. Gettierized Knobe Effects. *Episteme* 10 (3): 219–240.
- Blaine, B., and J. Crocker. 1993. Self-esteem and self-serving biases in reactions to positive and negative events: An integrative review. In *Self-esteem: The Puzzle of low Self-regard*, ed. R. Baumeister, 55–85. New York: Springer.
- BonJour, L. 2009. *Epistemology: Classic Problems and Contemporary Responses*. Lanham: Rowman & Littlefield Publishers.
- Buckwalter, W. 2012. Non-traditional factors in judgments about knowledge. *Philosophy Compass* 7 (4): 278–289.
- Buckwalter, W. 2014. Gettier made ESEE. *Philosophical Psychology* 27 (3): 368–383.
- Buckwalter, W., D. Rose, and J. Turri. 2015. Belief through thick and thin. *Noûs* 49 (4): 748–775.
- Byrne, A. 2005. Introspection. *Philosophical Topics* 33 (1): 79–104.
- Byrne, A. 2011. Knowing that I am thinking. In *Self-knowledge*, ed. A. Hatzimoysis, 105–124. Oxford: Oxford University Press.
- Cassam, Q. 1997. *Self and World*. Oxford: Oxford University Press.
- Cassam, Q. 2011. Knowing what you believe. In *Proceedings of the Aristotelian Society*, Vol. 111, part 1: 1–23.
- Cassam, Q. 2014. *Self-knowledge for Humans*. Oxford: Oxford University Press.
- Cialdini, R.B., S.L. Brown, B.P. Lewis, C. Luce, and S.L. Neuberg. 1997. Reinterpreting the empathy–altruism relationship: When one into one equals oneness. *Journal of Personality and Social Psychology* 73 (3): 481–494.
- Cohen, J. 2013. *Statistical Power Analysis for the Behavioral Sciences*. New York: Routledge.
- Cova, F. 2016. The folk concept of intentional action. In *A Companion to Experimental Philosophy*, eds. J. Sytsma and W. Buckwalter. Hoboken: John Wiley & Sons.
- Cova, F., and H. Naar. 2012. Side-effect effect without side effects: The pervasive impact of moral considerations on judgments of intentionality. *Philosophical Psychology* 25 (6): 837–854.
- Dalbauer, N., and A. Hergovich. 2013. Is what is worse more likely?—The probabilistic explanation of the epistemic side-effect effect. *Review of Philosophy and Psychology* 4 (4): 639–657.
- Dretske, F. 1994. Introspection. *Proceedings of the Aristotelian Society* 94: 263–278.
- Duff, A. 1980. Intention, Mens Rea and the law commission report. *Criminal Law Review*: 147–160.
- Duff, A. 1982. Intention, responsibility and double effect. *The Philosophical Quarterly* 32: 1–16.
- Duff, A. 1989. Intentions legal and philosophical. *Oxford Journal of Legal Studies* 9: 76–94.
- Epley, N., E. Caruso, and M.H. Bazerman. 2006. When perspective taking increases taking: Reactive egoism in social interaction. *Journal of Personality and Social Psychology* 91 (5): 872–889.
- Faul, F., E. Erdfelder, A. Buchner, and A.-G. Lang. 2009. Statistical power analyses using G* power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods* 41 (4): 1149–1160.

- Feldman, R. 2006. *Epistemology*. London: Pearson.
- Feltz, A. 2007. The Knobe effect: A brief overview. *Journal of Mind and Behavior* 28 (3/4): 265–277.
- Feltz, A., M. Harris, and A. Perez. 2012. Perspective in intentional action attribution. *Philosophical Psychology* 25 (5): 673–687.
- Fernandez-Duque, D., and T. Wifall. 2007. Actor/observer asymmetry in risky decision making. *Judgment and Decision making* 2 (1): 1–8.
- Galinsky, A.D., and G.B. Moskowitz. 2000. Perspective-taking: Decreasing stereotype expression, stereotype accessibility, and in-group favoritism. *Journal of Personality and Social Psychology* 78 (4): 708–724.
- Galinsky, A.D., G. Ku, and C.S. Wang. 2005. Perspective-taking and self-other overlap: Fostering social bonds and facilitating social coordination. *Group Processes & Intergroup Relations* 8 (2): 109–124.
- Galinsky, A.D., W.W. Maddux, D. Gilin, and J.B. White. 2008. Why it pays to get inside the head of your opponent: The differential effects of perspective taking and empathy in negotiations. *Psychological Science* 19 (4): 378–384.
- Gertler, B. 2010. *Self-knowledge*. London: Routledge.
- Hamman, G. 1976. Practical reasoning. *The Review of Metaphysics* 29 (3): 431–463.
- Holton, R. 2010. Norms and the Knobe effect. *Analysis* 70 (3): 417–424.
- Jones, E.E., and R.E. Nisbett. 1971. The actor and the observer: Divergent perceptions of the causes of behavior. In *Attribution: Perceiving the Causes of Behavior*, eds. E.E. Jones et al. Morristown: General Learning Press.
- Kauppinen, A. 2007. The rise and fall of experimental philosophy. *Philosophical Explorations* 10 (2): 95–118.
- Kneer, M., and S. Bourgeois-Gironde. 2017a. Attribution de mens rea: Données empiriques. In *Causalité, Responsabilité et Contribution à la Dette*, eds. S. Ferrey and F. G'Sell. Brussels: Éditions Bruylant. (Forthcoming).
- Kneer, M., and S. Bourgeois-Gironde, S. 2017b. Mens rea ascriptions, expertise and outcome effects: Professional judges surveyed. *Cognition* 169: 139–146.
- Knobe, J. 2003a. Intentional action and side effects in ordinary language. *Analysis* 63 (3): 190–194.
- Knobe, J. 2003b. Intentional action in folk psychology: An experimental investigation. *Philosophical Psychology* 16 (2): 309–324.
- Knobe, J. 2004a. Folk psychology and folk morality: Response to critics. *Journal of Theoretical and Philosophical Psychology* 24 (2): 270–279.
- Knobe, J. 2004b. Intention, intentional action and moral considerations. *Analysis* 64 (2): 181–187.
- Knobe, J., and A. Burra. 2006. The folk concepts of intention and intentional action: A cross-cultural study. *Journal of Cognition and Culture* 6 (1): 113–132.
- Knobe, J., and B.F. Malle. 2002. Self and other in the explanation of behavior: 30 Years later. *Psychologica Belgica* 42 (1/2): 113–130.
- Knobe, J., and G.S. Mendlow. 2004. The good, the bad and the blameworthy: Understanding the role of evaluative reasoning in folk psychology. *Journal of Theoretical and Philosophical Psychology* 24 (2): 252–258.
- Laurent, S.M., and M.W. Myers. 2011. I know you're me, but who am I? Perspective taking and seeing the other in the self. *Journal of Experimental Social Psychology* 47 (6): 1316–1319.
- Lehrer, K. 2015. *Theory of knowledge*. London: Routledge.
- 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 (5): 421–427.
- Lombrozo, T., and K. Uttich. 2010. Putting normativity in its proper place. *Behavioral and Brain Sciences* 33 (4): 344–345.
- Ludwig, K. 2007. The epistemology of thought experiments: First person versus third person approaches. *Midwest Studies in Philosophy* 31 (1): 128–159.
- Malle, B.F. 2006. The actor-observer asymmetry in attribution: A (surprising) meta-analysis. *Psychological Bulletin* 132 (6): 895–919.
- Malle, B.F., J.M. Knobe, and S.E. Nelson. 2007. Actor-observer asymmetries in explanations of behavior: New answers to an old question. *Journal of Personality and Social Psychology* 93 (4): 491–514.
- Maner, J.K., C.L. Luce, S.L. Neuberger, R.B. Cialdini, S. Brown, and B.J. Sagarin. 2002. The effects of perspective taking on motivations for helping: Still no evidence for altruism. *Personality and Social Psychology Bulletin* 28 (11): 1601–1610.
- McCullough, M.E., E.L. Worthington Jr., and K.C. Rachal. 1997. Interpersonal forgiving in close relationships. *Journal of Personality and Social Psychology* 73 (2): 321–336.
- Mele, A.R., and F. Cushman. 2007. Intentional action, folk judgments, and stories: Sorting things out. *Midwest Studies in Philosophy* 31 (1): 184–201.

- Mezulis, A.H., L.Y. Abramson, J.S. Hyde, and B.L. Hankin. 2004. Is there a universal positivity bias in attributions? A meta-analytic review of individual, developmental, and cultural differences in the self-serving attributional bias. *Psychological Bulletin* 130 (5): 711–747.
- Moore, M.S. 2011. Intention as a marker of moral culpability and legal punishability. In *Philosophical Foundations of Criminal Law*, eds. A. Duff and S.P. Green. Oxford: Oxford University Press.
- Murray, D., J. Sytsma, and J. Livengood. 2013. God knows (but does god believe?). *Philosophical Studies* 166 (1): 83–107.
- Myers-Schulz, B., and E. Schwitzgebel. 2013. Knowing that P without believing that P. *Noûs* 47 (2): 371–384.
- Nadelhoffer, T., and A. Feltz. 2008. The actor–observer bias and moral intuitions: Adding fuel to Sinnott-Armstrong’s fire. *Neuroethics* 1 (2): 133–144.
- Parker, S.K., and C.M. Axtell. 2001. Seeing another viewpoint: Antecedents and outcomes of employee perspective taking. *Academy of Management Journal* 44 (6): 1085–1100.
- Phelan, M.T., 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 (2): 291–298.
- Pollai, M., and E. Kirchler. 2012. Differences in risk-defusing behavior in deciding for oneself versus deciding for other people. *Acta Psychologica* 139 (1): 239–243.
- Robinson, B., P. Stey, and M. Alfano. 2015. Reversing the side-effect effect: The power of salient norms. *Philosophical Studies* 172 (1): 177–206.
- Rose, D., and J. Schaffer. 2013. Knowledge entails dispositional belief. *Philosophical Studies* 166 (1): 19–50.
- Schwitzgebel, E., and F. Cushman. 2012. Expertise in moral reasoning? Order effects on moral judgment in professional philosophers and non-philosophers. *Mind & Language* 27 (2): 135–153.
- Schwitzgebel, E., and F. Cushman. 2015. Philosophers’ biased judgments persist despite training, expertise and reflection. *Cognition* 141: 127–137.
- Siewert, C. 2012. On the phenomenology of introspection. In *Introspection and Consciousness*, eds. D. Smithies and D. Stoljar. Oxford: Oxford University Press.
- Smith, A. 1790/2002. *The Theory of Moral Sentiments*. Cambridge: Cambridge University Press.
- Sverdlik, S. 2004. Intentionality and moral judgments in commonsense thought about action. *Journal of Theoretical and Philosophical Psychology* 24 (2): 224–236.
- Takaku, S., B. Weiner, and K.-I. Ohbuchi. 2001. A cross-cultural examination of the effects of apology and perspective taking on forgiveness. *Journal of Language and Social Psychology* 20 (1–2): 144–166.
- Turri, J. 2014. The problem of ESEE knowledge. *Ergo* 1 (4): 101–127.
- Turri, J. 2016. Knowledge judgments in “Gettier” cases. In *A Companion to Experimental Philosophy*, eds. J. Sytsma and W. Buckwalter, 335–348. Hoboken: John Wiley & Sons.
- Uttich, K., and T. Lombrozo. 2010. Norms inform mental state ascriptions: A rational explanation for the side-effect effect. *Cognition* 116 (1): 87–100.
- Weijers, D. 2014. Nozick’s Experience machine is dead, long live the experience machine! *Philosophical Psychology* 27 (4): 513–535.
- Weinberg, J.M., J. Alexander, C. Gonneman, and S. Reuter. 2012. Restrictionism and reflection: Challenge deflected, or simply redirected? *The Monist* 95 (2): 200–222.
- Williamson, T. 2002. *Knowledge and its Limits*. Oxford: Oxford University Press.