

On Trying to Save the Simple View

THOMAS NADELHOFFER

Abstract: According to the analysis of intentional action that Michael Bratman has dubbed the ‘Simple View’, intending to x is necessary for intentionally x -ing. Despite the plausibility of this view, there is gathering empirical evidence that when people are presented with cases involving moral considerations, they are much more likely to judge that the action (or side effect) in question was brought about intentionally than they are to judge that the agent intended to do it. This suggests that at least as far as the ordinary concept of intentional action is concerned, an agent need not intend to x in order to x intentionally.

1. Setting the Stage

If you believe you were harmed intentionally, then you will likely be quite angry with the person who harmed you—especially if you feel the harm was unjustified. If, on the other hand, you believe you were harmed unintentionally, you will likely have an entirely different reaction. Owing to the important role that ascriptions of intentional action play in our daily lives, it is not surprising that the concept of intentional action often takes center stage in the action theory literature. But what does it mean to act intentionally and what is the relationship between intending to x and intentionally x -ing?*

According to the analysis of intentional action that Michael Bratman has dubbed the ‘Single Phenomenon View’ (SPV) (1987, p. 112), an agent who intentionally x -ed must have intended to do *something*—although she need not necessarily have specifically intended to x (1987, p. 113). Consider, for instance, the following example and comments from Alfred Mele:

Alice is mowing her lawn—intentionally, of course. In the process, she has taken many steps. It would be strange—and, I believe, mistaken—to maintain that her taking her next routine step while pushing the mower is not an intentional action. But there is no need to suppose that Alice has an intention specifically to take that step. Given that she intends to mow her lawn at the time, is a proficient mower, encounters no obstacles requiring alteration of her

I would like to thank Joshua Knobe, Hugh McCann, Al Mele, Eddy Nahmias, Virginia Tice, the editor of *Mind & Language*, and two anonymous referees for helpful comments on earlier versions of this paper.

Address for correspondence: Department of Philosophy, Dickinson College, Carlisle, PA 17013.

Email: tnadelhoffer@gmail.com

Mind & Language, Vol. 21 No. 5 November 2006, pp. 565–586.

© 2006 The Author

Journal compilation © 2006 Blackwell Publishing Ltd

gait, and so on, her mowing intention can do its work without her having a series of intentions corresponding to each routine step (Mele, 2005, p. 150; see, also, Mele, 1992, ch. 8).

On this view, Alice's general intention to mow the lawn may serve as a sort of proxy intention for taking all of the requisite steps for successfully mowing the lawn, steps that are intentionally taken despite the fact that Alice did not specifically intend to take each of them individually.

According to a competing analysis of intentional action—which Bratman has dubbed the 'Simple View' (SV)—in order for an agent to intentionally *x*, she must directly or specifically intend to *x* (Bratman, 1984, p. 377). On this view, Alice did not intentionally take each of the individual steps she took while cutting the grass unless taking precisely these steps was among the things that Alice intended to do at the time (e.g. Adams, 1986; McCann, 1986, 1991).¹ Whereas proponents of SPV claim that in order for an agent to intentionally *x* she must merely have acted upon some 'x-relevant intention',² proponents of SV state that the only x-relevant intention that will do the trick is the intention to *x* itself. So, if SV were a correct analysis of intentional action, then SPV would be correct as well—generally speaking—but the contrary is not the case. The truth of SPV does not entail the truth of SV.

Since SV is a particularized version of SPV—requiring a specific intention to *x* rather than merely a more general x-relevant intention—it would perhaps be most helpful to examine this view first. If nothing else, SV is a more demanding view of intentional action than SPV since it places more constraints on what counts as an intentional action. Proponents of this demanding yet plausible analysis of intentional action defend it on a number of grounds. First and foremost, it purportedly captures our pre-theoretical intuitions and coheres with our ordinary usage of the concepts of intending and intentional action (McCann, 1998, p. 210). According to proponents of SV, in ordinary contexts it would sound strange for us to say that someone intentionally *x-ed* even though she did not intend to *x*.

Second, given that SV is the etymologically intuitive view that *intending* to *x* is necessary for *intentionally x-ing*, SV has the virtue of being, well, simple or 'uncluttered' (Adams, 1986, p. 284). Third, SV 'gives us reason to believe that our

¹ While proponents of SV do claim that an intention to *x* is necessary for *x-ing* intentionally, they do not make the further claim that an agent who intends to *x* and actually does *x*, does *x* intentionally. On their view, the intention must cause the action "in the right way" for the action to be intentional (Adams, 1986, p. 284). By insisting that an agent's intentions must be connected to her actions in the right sort of way in order for those actions to be intentional, the proponent of SV thereby shield themselves from objections based on deviant causal chains whereby an agent intended to *x* and *x-ed*, but we would nevertheless not ordinarily say she *x-ed* intentionally.

² One of the referees for *Mind & Language* called these kinds of proxy intentions "x-relevant intentions"—a phrase that I found to be quite helpful. As far as I can tell, x-relevant intentions can come in at least two varieties: a) general intentions under whose umbrella particular intentional actions fall (see, e.g. Mele, 2005), and b) disjunctive intentions whereby an agent intends to do *x* or *y*, but not both, and succeeds (see, e.g. Bratman, 1987).

intentions causally guide our actions in virtue of their content' (Adams, 1986, p. 284)—thereby supporting our ordinary view of ourselves whereby the contents of our intentions to *x* play an important role in our intentionally *x-ing*. Finally, proponents of SV suggest that if intentions to *x* are *not* necessary for intentionally *x-ing*, in cases involving agents who lack an intention to *x*, we will be unable to distinguish intentionally *x-ing* from unintentionally *x-ing*.

Before we can judge the overall merits of either SV or SPV, we must first determine what it is precisely we are interested in when we are doing action theory. Whereas some philosophers focus primarily on developing explanatory models of the actual etiology of intentional actions, a number of philosophers and psychologists working under the rubric of 'experimental philosophy' have begun focusing instead on analyzing the concepts of ordinary language and investigating the intuitions of laypersons concerning intentions and intentional actions.³ My goal in this paper is not to suggest that one of these approaches is superior to the other. Indeed, on my view, both have their respective merits.

For present purposes, I will simply assume that understanding the relationship between the folk concepts of intention and intentional action is an important goal in its own right. Not only are appeals to common sense, ordinary intuitions, folk morality, folk psychology, ordinary usage, pre-theoretical judgments, etc. common in nearly all areas of analytic philosophy, but these appeals seem particularly pervasive and relevant in areas that often focus on our ordinary moral, political, and legal practices such as the free will debate, applied ethics, criminal theory, and action theory. Given the ubiquity of appeals to intuitions⁴ or to common sense—which are often straightforward empirical claims about ordinary usage—philosophers should not shy away from probing folk intuitions in an empirically informed way when these intuitions play an important role in the debate.⁵ Minimally, as more of

³ There are a growing number of philosophers in a variety of fields working under the rubric of 'experimental philosophy'. For a concise introduction to this developing trend, see Nichols (2004c). One of the primary goals of experimental philosophy is to increase our understanding of folk concepts and intuitions concerning issues that are especially relevant to philosophy. Hence, there is growing data about folk concepts and intuitions relevant to epistemology (e.g. Weinberg, Nichols and Stich, 2001; Nichols, Weinberg and Stich 2003), ethics (e.g. Doris and Stich, 2003; Nichols, 2004a), free will (e.g. Nahmias *et al.* 2005, forthcoming; Nichols, 2004b), and most importantly for our present purposes, the philosophy of action (Knobe, 2003a, 2003b, 2004; Knobe and Burra, forthcoming; Knobe and Mendlow, forthcoming; Malle and Knobe, 1997, 2001; Nadelhoffer, 2004a, 2004b, 2004c, 2005).

⁴ Determining precisely what intuitions are is a tricky matter that would take us too far afield for present purposes. For now, I am simply going to follow Alvin Goldman and Joel Pust's suggestion that, 'the contents of intuitions are usually singular classificational propositions, to the effect that such-and-such an example is or is not an instance of knowledge, of justice, of personal identity, and so forth' (1998, p. 182).

⁵ There is a growing literature addressing the meta-philosophical issues concerning the role that intuitions do and should play in philosophy. DePaul and Ramsey (1998) contains several interesting and insightful articles on intuitions in philosophy and psychology. More recently, Brian Weatherson (2003), Ernest Sosa (2005), Michael Lynch (2005), and Timothy Williamson (2004) have attempted to flesh out the relevance of intuitions for philosophy.

these preliminary steps are taken, we will slowly pave the way towards a richer and more robust understanding of the relationship between folk morality and folk psychology—particularly the role played by foresight, intentions, and intentional actions in judgments concerning moral and legal responsibility.

As our knowledge about folk intuitions and concepts continues to grow, it will become increasingly difficult for philosophers to claim that their position aligns with common sense unless their views empirically merit such support. Of course, if a particular analysis of an ordinary concept *does* turn out to cohere with folk intuitions that alone would not prove it to be *true*, but it would seem to shift the argumentative burden to those who argue *contrary* to these intuitions.⁶ By my lights, the analysis of intentional action that most closely agrees with the judgments of non-specialists enjoys ‘squatters’ rights’ until it is shown to be defective for other reasons.⁷ Many parties to the debate have simply assumed that SV enjoys precisely this honor without checking to see whether this assumption is supported by anything more than anecdotal evidence.

As Mele has correctly pointed out, any adequate philosophical analysis of intentional action should minimally be ‘anchored by common-sense judgments’ about particular cases (Mele, 2001, p. 27)—even if it need not capture or reflect *all* of these judgments. And the only method of determining what the majority of non-specialists say about particular cases is to *actually ask them*. Having done so, if we find that an analysis of intentional action is entirely inconsistent with folk intuitions, we will be in a good position to suggest that it ‘runs the risk of having nothing more than a philosophical fiction as its subject matter’ (Mele, 2001, p. 27). Data about the folk concept of intentional action become all the more important for proponents of SV such as Hugh McCann who explicitly claim to be interested in ordinary concepts and not their philosophical counterparts (McCann, 1998, p. 210). Minimally, any philosopher who offers an account of intentional action that is not anchored by folk judgments would need to admit that her view does not cohere with some aspects of ordinary usage.⁸

In a groundbreaking paper on the folk concept of intentional action, Bertram Malle and Joshua Knobe made one of the first attempts to examine this concept in an empirically informed way (Malle and Knobe, 1997). In one study that is

⁶ See Weatherson (2003), for a discussion of the role played by intuitions in epistemology—especially concerning Gettier cases.

⁷ It is worth pointing out that the philosophical relevance of folk intuitions will vary from topic to topic. Owing to the inherently abstract nature of some philosophical subjects—e.g. logic or the philosophy of mathematics—we may not care at all about pre-theoretical judgments and intuitions. In other areas of philosophy, however, it looks as if folk intuitions do (and arguably should) play a more important role in philosophical theories—especially when the relevant concepts are intimately bound up with our everyday moral views and practices.


⁸ Neither I nor anyone else that I am aware of suggests that we can legitimately move from ‘intuitive’ to ‘true’. The fact that a particular analysis or philosophical position is widely intuitive does not make it any more likely to be true, just much harder for people to abandon in the event that it turns out to be fatally flawed.

particularly germane to our present discussion, 159 undergraduates were given a questionnaire and asked to answer the following question: ‘When you say that somebody performed an action intentionally, what does that mean? Please explain’ (Malle and Knobe, 1997, p. 8). All of the answers were then transcribed into another booklet and then coded. As part of the coding process, synonyms for the term ‘intentionally’ (e.g. ‘purposely’ or ‘deliberately’) were not counted. Once these cognates were excluded from the data set, the remaining answers revealed that there were apparently four main components: desire, belief, intention, and awareness.

Because these four basic components accounted for 96% of the participants’ definitions, Malle and Knobe conclude that ‘the folk concept of intentional action includes: (a) a desire for an outcome, (b) beliefs about the action leading to that outcome, (c) an intention to perform that act, and (d) awareness of performing that act’ (Malle and Knobe, 1997, pp. 8-9). Malle and Knobe also suggest that the intention component serves as a link between an agent’s desires and beliefs and her actions. As they say, ‘people appreciate that an intentional action does not derive from desire and belief alone but that its direct cause is an intention’ (Malle and Knobe 1997, p. 9).

As Malle and Knobe point out, none of the past psychological models of intentional action capture all of the elements of their own four-component model. Whereas Jones and Davis (1965), Ossorio and Davis (1968), and Shaver (1995) all identified the belief and desire components, they overlooked the importance of intention and awareness. Similarly, Heider’s (1958) model of intentional action had both an intention and a desire component, but it lacked belief and awareness components. Nevertheless, Malle and Knobe point out that all of these models also postulated something that their own four-component model lacks—namely, an ability or skill component. Surprisingly, this component was absent from the participants’ definitions—which explains why Malle and Knobe do not include it in their model. But given that it is intuitively plausible that ability or skill are necessary for intentional action, Malle and Knobe ran a few additional studies which did in fact show that folk ascriptions of intentional action are sensitive to considerations of skill and control after all (Malle and Knobe, 1997, pp. 10-12; see, also, Knobe, 2003b; Nadelhoffer, 2005).

Consequently, Malle and Knobe amended their original model of the folk concept of intentional action to include a skill component. The resulting five-component model looks like this:



According to this revised model, performing an action intentionally ‘requires the presence of five components: a desire for an outcome; beliefs about an action that leads to that outcome; an intention to perform the action; skill to perform the action; and awareness of fulfilling the intention while performing the action’ (Malle and Knobe, 1997, p. 12). Moreover, Malle and Knobe suggest that these five components are ‘hierarchically arranged, such that belief and desire are necessary conditions for attributions of intention and, given an *intention*, skill and awareness are necessary conditions for attributions of *intentionality*’ (Malle and Knobe, 1997, p. 15).

Malle and Knobe’s five-component model of the folk concept of intentional action has several strengths. First, it is empirically informed rather than merely speculative. Second, by purportedly identifying all of the necessary components of the folk concept of intentional action, Malle and Knobe’s model not only integrates a number of the past analyses of intentional action that have been put forward by both psychologists and philosophers, but it also reveals precisely where these previous analyses went wrong. Third, it corrects mistakes that are found in previous models—especially concerning the proper way to understand the relationship between skill, control, and folk ascriptions of intentional action. Finally, their model helps to ‘clarify a nagging terminological complexity’—namely, the relationship between intentions and intentional actions (Malle and Knobe, 1997, p. 16).

The most important result of Malle and Knobe’s research for present purposes is that it appears on the surface to support SV as an analysis of the folk concept of intentional action. After all, according to their five-component model, people do not ordinarily say that an agent intentionally *x-ed* unless she intended to *x*. However, subsequent research has shown that their model fails to adequately capture the full range of folk intuitions concerning intentions and intentional actions. More specifically, we now know that Malle and Knobe’s earlier studies overlooked the possibility that moral considerations may have a striking effect on folk ascriptions of intentional action. Thanks to recent work being done by Knobe and others, we have learned that when participants are given vignettes that involve morally laden actions—especially bad ones—they often ascribe intentionality⁹ to these actions even though they do *not* judge that the agent intended to perform these actions. As we are about to see, the results of several recent studies cast doubt on the claim that SV coheres with our ordinary intuitions concerning intentions and intentional action.

2. Knobe’s CEO Studies

Following up on his earlier work with Malle, Knobe recently set out to determine whether folk intuitions about the intentionality of foreseeable yet undesired side

⁹ For the purposes of this paper, whenever I discuss intentionality, I am only talking about the question of whether an agent’s actions are intentional. This sort of intentionality is to be distinguished from discussions of intentionality that one finds in the literature on the philosophy of mind. When philosophers talk about intentionality in this latter context, they are usually interested in the question of how some of our mental states can be *about* things in the world.

effects are influenced by moral considerations (Knobe, 2003a). Each of the 78 participants in the first of Knobe's side effect studies were presented with a vignette involving either a 'harm condition' or a 'help condition'. Those who received the harm condition read the following vignette:

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

They were then asked to judge how much blame the chairman deserved for harming the environment (on a scale from 0 to 6) and to say whether they thought the chairman harmed the environment intentionally. 82% of the participants claimed that the chairman harmed the environment intentionally. Participants in the help condition, on the other hand, read the same scenario except that the word 'harm' was replaced by the word 'help'. They were then asked to judge how much praise the chairman deserved for helping the environment (on a scale from 0 to 6) and to say whether they thought the chairman helped the environment intentionally. Only 23% of the participants claimed that the chairman intentionally helped the environment (Knobe 2003a, p. 192).

In another side-effect study, Knobe got similar results.¹⁰ When Knobe combined the praise and blame ratings from the two studies, he got the following results: Whereas the participants who were given the harm condition said the agent deserved a lot of blame ($M=4.8$), those who were given the help condition said

¹⁰ This time each of the 42 participants received one of the following two vignettes: *Harm Condition*: A lieutenant was talking with a sergeant. The lieutenant gave the order: 'Send your squad to the top of Thompson Hill'. The sergeant said: 'But if I send my squad to the top of Thompson Hill, we'll be moving the men into the enemy's line of fire. Some of them will surely be killed!' The lieutenant answered: 'Look, I know that they'll be in the line of fire, and I know that some of them will be killed. But I don't care at all about what happens to our soldiers. All I care about is taking control of Thompson Hill'. The squad was sent to the top of Thompson Hill. As expected, the soldiers were moved into the enemy's line of fire, and some of them were killed. *Help condition*: A lieutenant was talking with a sergeant. The lieutenant gave the order: 'Send your squad to the top of Thompson Hill'. The sergeant said: 'But if I send my squad to the top of Thompson Hill, we'll be taking them out of the enemy's line of fire. They'll be rescued!'. The lieutenant answered: 'Look, I know that we'll be taking them out of the line of fire, and I know that some of them would have been killed otherwise. But I don't care at all about what happens to our soldiers. All I care about is taking control of Thompson Hill'. The squad was sent to the top of Thompson Hill. As expected, the soldiers were moved out of the enemy's line of fire, and some of them were saved. Once again, the harm and the help conditions yielded drastically different responses: 77% of the participants who read the harm condition said the agent intentionally brought about the negative side effect, whereas only 30% of those who read the help condition said the agent brought about the positive side effect intentionally (Knobe, 2003a, p. 192).

that the agent deserved virtually no praise ($M=1.4$). Moreover, these results were correlated with their judgments about whether or not the side effect was brought about intentionally (Knobe, 2003a, p. 193).

One of the most noteworthy features of Knobe's side effect studies is that in the harm conditions, participants judged that the side effects were brought about intentionally even though neither the CEO nor the lieutenant cared at all about doing so. If we assume for the sake of argument that an agent does not intend to x if she does not care at all about whether she x -es, then, at first blush at least, Knobe has generated empirical evidence which suggests that SV is not supported by folk intuitions. After all, an overwhelming majority of participants in the harm condition of Knobe's CEO study judged that the CEO intentionally harmed the environment even though he did not care, and hence did not intend, to do so. Of course, not everyone agrees with this interpretation of the data.

3. The Pragmatic Features of Intentional Language vs. The Core Concept

Fred Adams and Annie Steadman have suggested that Knobe's studies get at the 'pragmatic dimension of intentional talk' rather than getting at what they call the 'core concept' of intentional action (2004a, p. 174). On their view, understanding the concept of intentional action requires both an understanding of the 'cognitive machinery that underlies intentional action' and a grasp of the type of counterfactual dependency that necessarily holds between intentions and intentional action (Adams and Steadman, 2004a, p. 174). And to the extent that laypersons purportedly lack an understanding of these underlying concepts and causal relationships, 'folk notions of intentional action are not clearly articulated' (Adams and Steadman, 2004a, p. 173). Hence, Adams and Steadman conclude that rather than accessing a core concept of intentional action, Knobe's studies are merely getting at participants' grasp of the pragmatic features of intentional language—features that include the types of judgments that people make that depend more on social context than on the 'semantic content of a sentence or judgment' (2004a, p. 174).

Adams and Steadman claim that to say that an agent is blameworthy for x -ing is to conversationally imply that she intentionally x -ed—even though blame is purportedly not part of the semantic content of intentional language. To support this latter claim, they suggest that, 'the truth conditions for 'S did A intentionally' do not include praise or blame. It is not necessary for act A to be good or bad for the action to be intentional' (Adams and Steadman, 2004a, p. 178). Hence, according to Adams and Steadman, this sort of conversational implicature—and not an underlying folk concept of intentional action—must have been driving the participants' judgments in Knobe's studies. So, for instance, because the participants disapproved of the CEO's indifference to harming the environment, they wanted to blame him. Moreover, since to say that the CEO is blameworthy for harming the environment is to conversationally imply that he intentionally did so, the

participants judged that his harming the environment was intentional. But given that blame does not constitute part of the semantic meaning of the core concept of intentional action, the results of Knobe's experiment purportedly only give us insight into the pragmatic features of intentional language.

In response to Adams and Steadman, Knobe conducted a new study that was designed to eliminate the effect of pragmatic implicature—thereby undermining the force of their objection to his view (Knobe, 2004). In devising this new study, Knobe used the very same harm and help CEO vignettes from his earlier side effect studies, except this time he substituted 'in order to' for 'intentionally' in the questions the participants were asked to answer. In revising the old CEO vignettes in this way, Knobe assumes that 'people's use of the phrase 'in order to' thereby provides us with a kind of indirect evidence about which behaviors they regard as intentional' (Knobe, 2004, p. 183).

Each of the 77 participants was randomly assigned to either the harm condition or the help condition and then given the relevant version of the following sentence: 'The chairman harmed [helped] the environment in order to increase profits'. Participants were then asked to judge whether this sentence 'sounded right to them' (Knobe, 2004, p. 184) on a scale from -3 ('sounds wrong') to +3 ('sounds right'). The results were as follows: the average rating for participants in the harm condition was +.6 and the average rating in the help condition was -1. The difference between the two conditions was statistically significant [$t(77) = 2.65$, $p = 0.01$].

According to Knobe, by (a) framing the questions in terms of 'in order to' rather than 'intentionally', and (b) asking participants to judge whether a particular English sentence sounded right to them rather than asking them whether side effects of the CEO's actions were brought about intentionally, he has avoided Adams and Steadman's concern that the pragmatic rather than the semantic features of intentional language best explain the results of his earlier studies. Hence, he concludes that, 'it seems unlikely that the difference between people's responses to the harm vignette and their responses to the help vignette is due entirely to pragmatic features. At this point, the most plausible hypothesis seems to be that the difference between the two vignettes is showing something fundamental about people's concept of intentional action' (Knobe, 2004, p. 184)—a conclusion that Adams and Steadman have subsequently denied (Adams and Steadman, 2004b).

In response to Knobe's attempt to get around their earlier objection, Adams and Steadman have subsequently pointed out that his new study falls prey to the very same line of reasoning. After all, on their view, stating that an agent did x in order to do y is just another way of stating that she did y intentionally (Adams and Steadman, 2004b, p. 270). As they say, 'consider the sentence 'S did A intentionally'. Clearly this is intentional talk. It presupposes a purpose or goal that constitutes the reason why S did A. S might have done A in order to B. Or S might have done A for its own sake, but, being intentional, there is some reason why S did A' (Adams and Steadman, 2004b, p. 271). In this respect, they are in complete agreement with Knobe's claim that reasons-explanations are only applicable to

intentional actions. Yet, this is precisely why Adams and Steadman deny that Knobe has successfully gotten around their earlier objection. After all, if all intentional actions have corresponding reasons explanations, then ‘Knobe’s new wording still is intentional talk, and it’s ripe for pragmatic features’ (Adams and Steadman, 2004b, p. 270).

4. Intentions and Intentional Action: A Follow-Up Study

Now I want to consider another objection to Knobe’s work that has been put forward by Adams and Steadman—namely, that nothing in Knobe’s first side effect studies justifies the assumption that participants were not judging that the CEO intended to harm the environment (Adams and Steadman, 2004a). After all, Knobe never specifically asked the participants about the CEO’s intentions. And while it is plausible that because the CEO did not care at all about harming the environment the participants did not judge that he intended to do so, Adams and Steadman correctly point out that ‘it is *at least possible* that in the minds of the folk, the actors *did intend* the respective outcomes’ (2004a, p. 180). If so, then the results of Knobe’s earlier studies would not be inconsistent with SV after all.

To address this second objection, Knobe conducted yet another study (Knobe, 2004) in an effort to determine whether participants judge that the CEO had the intention of harming the environment. This time, the 63 participants were assigned randomly to either a harm or a help condition, and within each of these conditions participants were further divided into an ‘intentionally’ condition and an ‘intention’ condition. Those in the ‘intentionally’ condition were asked whether the CEO harmed (or helped) the environment intentionally and those in the ‘intention’ condition were asked whether it was the CEO’s intention to harm (or help) the environment. The results were as follows:

ANSWERS:	Harm condition	Help condition
Intentionally	87%	20%
Intention	29%	0%

As Knobe points out, ‘the most striking result...was that relatively few people said that it was the chairman’s intention to harm the environment. Within the harm conditions, we therefore obtain a significant difference between people’s responses for ‘intention’ and their responses for ‘intentionally’, $\chi^2(1, N = 32) = 10.6, p < 0.01$ (p. 185). From these results, he concludes that he has ‘identified a behavior such that (1) most people don’t think that the agent had an intention to perform it, but (2) most people do think that the agent performed it intentionally’ (Knobe, 2004, p. 186). Thus, he believes that he has successfully generated the empirical evidence needed to undermine the truth of SV. After all, a majority of

the participants in the harm condition were much more likely to judge that the CEO intentionally harmed the environment than they were to judge that the CEO acted with the intention of doing so. On the surface, these results suggest that Knobe has provided persuasive evidence that intending to x is not necessary for intentionally x -ing—at least as far as the folk concept is concerned.

5. The Objection from Cognitive Dissonance

In responding to Knobe's latest version of the CEO studies, Adams and Steadman claim that because Knobe did not ask each participant *both* the intentionality question *and* the intention question, the results of his new studies still fail to falsify SV (Adams and Steadman, 2004b). By their lights, 'if Knobe were to put on his questionnaire 'Is it possible to A intentionally without intending to do A?' the folk would experience cognitive dissonance. They would likely hesitate to embrace both that the chairman intentionally harmed the environment and that he did not intend to harm the environment' (Adams and Steadman, 2004b, p. 276). Hence, because Knobe did not force the participants to explicitly assert something that is inconsistent with SV, Adams and Steadman claim that he has yet to generate data against the truth of the claim that intending to x is necessary for intentionally x -ing. This objection to Knobe has also been put forward recently by Hugh McCann (forthcoming).

McCann has recently acknowledged that when we take the results of Knobe's studies at face value, 'they appear to sound the knell of SV, at least as it is applied to actions with wrongful side effects recognized in advance by the agent' (forthcoming). After all, a majority of the respondents judged that the CEO harmed the environment intentionally whereas only a minority of them judged that the CEO acted with a corresponding intention to do so. Nevertheless, McCann suggests that:

Defenders of SV have reason, however, to call for further evidence. An interesting fact about the second of the above studies is that none of the subjects was asked to pronounce on *both* the issue of intentionality and that of intention. That is, none were tested on their willingness to make the pair of statements that, according to SV, would be implicitly contradictory: namely, that the chair in either vignette had acted intentionally yet lacked the corresponding intention (McCann, forthcoming).

McCann is here making the very same point expressed earlier by Adams and Steadman—namely, that if participants were forced to explicitly deny SV, they would experience a cognitive dissonance that would lead them to say either (a) that the CEO both intentionally harmed the environment and had an intention to do so, or (b) that the CEO did not intentionally harm the environment and did not have an intention to do so. Given that this, too, is a straightforward empirical

prediction about how participants would respond to a given scenario, McCann correctly sets out to test it with a new study that Knobe helped him design.

McCann's study is also based on Knobe's earlier CEO harm vignette, except this time participants are asked whether the CEO harmed the environment intentionally and also one of the following two questions: (a) whether the chairman *had the intention* to harm the environment, or (b) whether the chairman *intended* to harm the environment. The motivation for asking some participants whether the CEO intended to harm the environment and others whether he acted with the intention of doing so is based on McCann's suspicion that to judge that an agent acted with the intention of *x-ing* is to judge that *x-ing* was the agent's *primary* goal. In this respect, McCann was trying to correct for the possibility that in Knobe's earlier study, the only reason participants did not judge that the CEO acted with the intention of harming the environment is that they correctly judged that harming the environment was not the CEO's main motivation or goal.

Keeping these distinctions in mind, McCann ran two studies—one with a 'have an intention' condition and one with an 'intended' condition. In the first study—which contained the 'have an intention' condition—the 106 participants were divided at random into four groups—those in the first two groups were asked only one of the questions whereas those in the second group were asked both questions. The results were as follows:

ANSWERS	One question	Both questions
Intentionally	63%	80%
Intention	27%	12%

As you can see, the results displayed in the left column are roughly in line with the results of Knobe's earlier study. The results displayed in the right column, on the other hand, seemingly falsify McCann's prediction that when participants are faced with having to say something inconsistent with SV, cognitive dissonance would minimize the asymmetry in their judgments concerning whether the CEO intentionally harmed the environment and whether he acted with the intention of doing so. Indeed, it had the opposite effect! As McCann points out, 'to the further dismay of SV, the contrast becomes even more pronounced when both questions are asked of the same subjects' (forthcoming).

In the second study—where the phrase 'have an intention' was replaced with the term 'intended'—there were 99 participants, divided once again into four groups, parallel to those in the first study. The results were as follows:

ANSWERS	One question	Both questions
Intentionally	64%	75%
Intended	42%	31%

As McCann had predicted, the change in wording made a difference—especially for those participants who received both questions. In the first study, only 12% of the participants who received both questions judged that the chairman had an intention to harm the environment, whereas 31% of those who received both questions in the second study judged that the chairman intended to harm the environment. McCann's worry about the pragmatic implicature of the phrase 'have an intention' was apparently justified. However, even once he corrected for this problem with the new wording, there was *still* a highly statistically significant difference in the second study between participants' judgments of intentionality and their judgments of intention—especially when participants were given both questions. In response to these results, McCann suggests that:

Indeed, they [i.e. the participants] seemed to welcome the opportunity to understand the case in just this way, with those ascribing intentionality to the chairman increasing, and those imputing intending/intention decreasing. This appears to be the precise opposite of shying away from any contradiction implicit in such a pair of verdicts, and so may be thought to finish off SV once and for all (McCann, forthcoming).

McCann's response to the results of the second study is even more surprising than the results themselves: he suggests that, 'perversely enough, it is in fact SV that is behind the phenomenon' (forthcoming).

Keep in mind that an overwhelming majority of the participants were willing to judge that the CEO harmed the environment intentionally even though he did not intend to do so—thereby embracing the alleged contradiction that both McCann and Adams and Steadman had predicted would generate enough cognitive dissonance to minimize the asymmetry in participants' judgments. Given that their prediction failed, McCann is correct to worry that he might look like a 'bad sport' for suggesting that the very same SV that is seemingly called into question by his study actually helps explain the results. Ultimately, McCann simply follows Adams and Steadman in suggesting that the pragmatic moral features of intentional language coupled with the truth of SV are responsible for the participants' responses and not the semantic meaning of 'intentional'. As he says:

If this is right, then the factors that appear to tip the results of our study against SV may for the most part be pragmatic rather than anything to do with the semantics of the terms 'intentional' and 'intended'. Moreover, these pragmatic features would not obtain but for the fact that when ordinary speakers address cases like the harm vignette, SV is one principle at work in their thinking (McCann, forthcoming).

By my lights, McCann's response works no better than Adams and Steadman's earlier objection to Knobe's studies. To see why, I ran a few studies of my own.

But before discussing the results of these studies in §7, I would first like to show that McCann has failed to properly appreciate the relevance of Knobe's data to the question of the truth of the claim that SV accurately describes the folk concept of intentional action.

6. The Problem of Non-Falsifiability

Given that McCann, unlike Adams and Steadman, has explicitly claimed to be interested in the *ordinary* concept of intentional action rather than one of its technical philosophical counterparts (McCann, 1998, p. 210), it is particularly important that his version of SV is consistent with the gathering data about ordinary usage. And perhaps the best way to determine whether SV coheres with folk intuitions is to treat it as a prediction about how people would ordinarily respond to certain kinds of scenarios. Insofar as SV is the claim that in order for an agent to intentionally *x*, she must have intended to *x*, we can safely assume that proponents of SV who claim that it is line with ordinary usage are committed to the following prediction: If you give laypersons a case involving an agent who does *x* and they judge that the agent *x-ed* intentionally, then they will also judge that the agent intended to *x*. Indeed, this is precisely what McCann had in mind when he revised Knobe's original studies. And, much to his surprise, the participants were willing to say that the CEO intentionally harmed the environment even though he did not intend to do so.

As McCann correctly points out, these results seemingly sound the knell of SV—treated here as a claim about the ordinary concept of intentional action. After all, since SV is the claim that intending to *x* is necessary for intentionally *x-ing*, what better evidence could McCann have for the falsity of SV than the fact that an overwhelming majority of participants in his own studies regarded an agent as *x-ing* intentionally whom they also regarded as not intending to *x*. At some point, it begins to look like both McCann and Adams and Steadman are simply making *ad hoc* appeals to the pragmatic features of intentional language in an effort to save SV from the empirical data at all costs.

Considering McCann's aforementioned interest in capturing the ordinary concept of intentional action, it is fair of us to ask him the following question: If the results of your own studies do not falsify SV as an analysis of the ordinary concept of intentional action, then what would? If the answer is nothing at all—to the extent that all data that conflict with SV will simply be explained away along Gricean lines—then McCann can no longer claim to be interested in the ordinary concept. For, how else are we supposed to determine what this concept is if not by surveying laypersons to find out how they apply it in particular cases?

If, on the other hand, McCann has a clear idea of what kind of evidence *would* count against SV, then he should explicitly spell it out. Because as far as I can tell, the best possible evidence that we can have that SV does not accurately describe the folk concept of intentional action is that folk ascriptions of intentional action

are inconsistent with predictions entailed by SV. We cannot have it both ways. We cannot insist that we are interested in the folk concept of intentional action rather than one of its technical counterparts while at the same time rejecting the results of the only kind of studies that might enable us to get at this concept.

Perhaps McCann has not spelled out what he would take as evidence against SV because he believes his opponents have misunderstood the significance of empirical data in the first place. After all, as he says, 'it is, first of all, a mistake to treat studies like these as a kind of election: that is, to declare that those who describe the chairman's actions as intentional but not intended, being in the majority, must represent "the" ordinary conception of intention' (McCann, forthcoming). McCann makes these remarks when discussing the results of his own studies. More specifically, he is responding to the possible suggestion that his second study militates against the truth of SV—something McCann denies.

The first thing worth pointing out, is that when McCann makes this claim about the inappropriateness of applying majority rules to the results of these kinds of studies, he appeals to the 42% of the participants in the first group of study two who said that the CEO intended to *x*. However, given that McCann explicitly claims that the results of the second group—where participants were forced to judge whether the CEO intentionally harmed the environment *and* whether he intended to do so—are more central to the debate about SV, he should use the figures from the second group rather than those from the first. And since McCann was the one who insisted on the importance of asking the questions together in the first place, he cannot fall back on the data from the study where only one question was asked.

More importantly, McCann seems to miss the point of the empirical objection to SV. As we have already seen, he claims that according to ordinary usage, intending to *x* is a necessary condition for intentionally *x-ing*. But if more than half of all of the subjects do not use the concept in this way, then I see no way of saving SV so conceived. After all, if the majority of folk intuitions about particular cases are inconsistent with SV—which they generally are—and if it does not appear that a number of the participants are 'conceptually misguided' by the language of the vignettes—which it does not appear that they are—then I am unsure what else would show that SV fails as an analysis of the ordinary concept of intentional action.

Of course, one can try to show that even those participants whose judgments are not consistent with SV actually rely on SV in their deliberations—much as McCann himself tries to do—but that does not undermine the fact that if we assume for the sake of argument that 97% of all subjects judged that the CEO *x-ed* intentionally but did not intend to *x*, this would indeed spell trouble for proponents of SV. This possibility notwithstanding, he claims that proponents of SV should not feel obligated to defend the view against 'any and all challenges', nor need they worry about any and all exceptions to the principle that may be revealed by a close examination of the relevant empirical data. However, given that the data strongly support the claim that *SV is the exception rather than the rule* when it comes to ascriptions of intention and intentional action in ordinary language, it is unclear

that proponents of SV can simply dismiss the counter-evidence on the grounds that it is uncommon or extraordinary without at the same time undermining their ability to maintain that SV accords with ordinary usage.

Minimally, proponents of SV owe us an account of what they mean when they say that according to ordinary language intending to *x* is necessary for intentionally *x-ing*. They also need to specify what kind of evidence would count as falsifying SV—at least as far as the folk concept of intentional action is concerned. I have suggested that claims about ordinary language must be sensitive to the relevant data. More specifically, I have claimed that the easiest way of holding the feet of SV to the empirical fire is to treat it as a prediction about how ordinary people will respond to specific scenarios. The results of the new studies I am about to discuss give us additional reason to doubt whether proponents of SV can continue to maintain in good faith that their view best captures our ordinary intuitions concerning intentional action.

7. A New Study

As we have already seen, both McCann and Adams and Steadman try to get a lot of mileage out of the distinction between the pragmatic features of intentional language, on the one hand, and the semantic content of the concept of intentional action, on the other hand. More specifically, they explain Knobe's data away by suggesting that the negative moral features of Knobe's CEO harm vignettes are doing all of the work, not the core concept itself. If it could be shown that people are willing to judge that a morally *neutral* side effect was brought about intentionally even though the agent did not intend to bring it about, then this would effectively preclude proponents of SV from being able to appeal to the pragmatic moral features of intentional language in their attempts to rescue SV from the empirical data. So, I conducted a series of studies.

Participants were forty undergraduates—each of whom read the following vignettes inspired by Gilbert Harman (1978):¹¹

Case 1:

A sniper has been ordered to kill an enemy commander. So, after getting himself into position, he finally has the enemy commander in his sights. Before he pulls the trigger, however, the sniper realizes that the gunfire will definitely alert the other enemy soldiers to his presence. But the sniper doesn't care at all about that—he just wants to shoot his target. So, he pulls the trigger—thereby shooting and killing the commander. And, as he expected, the enemy soldiers are alerted to his presence.

¹¹ I ran a number of other similar studies in an effort to probe folk intuitions concerning degrees of foresight. For the full study, see Nadelhoffer (forthcoming b).

Each participant was then asked the following questions: (1) Did the sniper intentionally alert the enemies to his presence? (2) Did the sniper want to alert the enemies to his presence? (3) Did the sniper intend to alert the enemies to his presence? (4) Did the sniper purposely alert the enemies to his presence? The results were as follows:

ANSWERS	C1
Intentionally	55%
Wanted	10%
Intended	10%
Purposely	55%

Given that the side effect in this case is *morally neutral*, it should not be surprising that the intentionality rating in these studies was not nearly as high as it was in Knobe's CEO harm cases. Nevertheless, there is still a highly statistically significant difference between the answers to questions (1) and (3)—($\chi^2(1, N=40) = 18.46, p < 0.001$). And since the proponents of SV cannot explain these results away in terms of the pragmatic moral features of intentional language, we now have further reason to doubt whether SV accurately describes the folk concept of intentional action.¹²

However, to the extent that the sniper nevertheless had a prudential reason *not* to alert the enemies to his presence, perhaps proponents of SV will simply revise their pragmatic implicature objection to explain why participants in my sniper cases were more willing to say that the sniper intentionally alerted the enemies to his presence than they were to say that he intended to do so.¹³ To head off this objection, I ran another study. This time, participants were 40 undergraduates, each of whom received the following vignette:

Before he pulls the trigger, however, the sniper realizes that the gunfire will definitely cause the barrel of his gun to get hot. But the sniper doesn't care at all whether the barrel of the gun is hot, he doesn't have to touch it anyway.

¹² In one respect, the results of C2 and C3 appear to be consistent with SV insofar as a majority of participants judged that the sniper neither intentionally nor intended to alert his enemies. Nevertheless, there is still a statistically significant difference between the number of participants who judged that the sniper intentionally alerted his enemies—20% and 40% respectively—and the number of those who judged that he intended to do so—0% and 15% respectively. And given that this is precisely what both McCann and Adams and Steadman predicted would not occur if participants were asked both questions, this is more problematic for the SV than it may seem at first blush.

¹³ We find this line of reasoning in Harman (1976). According to Harman's 'reasons against' view, if an agent (a) foresees that by doing *x* she will bring about *y*, (b) she has a reason not to bring *y* about, and (c) she does *x* intentionally, then she brings about *y* intentionally as well, even if she neither wanted, aimed, nor intended to bring *y* about.

So, he pulls the trigger—thereby shooting and killing the commander. And, as the sniper expected, firing the gun caused the barrel to heat up.

Participants were then asked the following four questions: (1) Did the sniper intentionally heat up the barrel of his gun? (2) Did the sniper want to heat up the barrel of his gun? (3) Did the sniper intend to heat up the barrel of his gun? (4) Did the sniper purposely heat up the barrel of his gun? The results were as follows:

ANSWERS	C2
Intentionally	68%
Wanted	5%
Intended	28%
Purposely	60%

When we once again compare the participants' responses to the first and third questions, there is a highly statistically significant difference between their judgments concerning whether the sniper intended to heat up the barrel of his gun and whether he intentionally did so—(χ^2 (1, N=40) = 12.84, $p < 0.001$).

Given that (a) the side effect of this vignette is morally neutral, and (b) the sniper has no reason not to bring about the side effect, the results of this study not only further undermine the claim that SV settles with ordinary usage, but they are not susceptible to the charge that pragmatic features of intentional action linked to negative moral or prudential considerations—and not the semantic meaning of the folk concept 'intentional'—explain the participants' judgments. Minimally, the results of these new surveys cast further doubt upon the claim that SV accurately captures the full range of our ordinary intuitions. But if even this sort of evidence is deemed to be insufficient for falsifying SV, then proponents of SV should spell out exactly what kind of evidence would do the trick.

8. The Single-Phenomenon View Revisited

Before closing, it is worth pointing out that even though the lion's share of the data concerning the folk concept of intentional action is *inconsistent* with SV—or so I have argued—these data are nevertheless *consistent* with SPV—i.e. the view that in order to intentionally x , and agent must act upon some x -relevant or proxy intention, even if she need not necessarily have specifically intended to x . In C1, for instance, it seems clear that the sniper did not intend to alert his enemies to his presence any more than he intended to make a loud noise by firing his gun. After all, if he had had a silencer at his disposal, he would certainly have used it since

doing so would have allowed him to kill the enemy commander, something he intended to do, without at the same time alerting his enemies to his whereabouts, something the sniper neither wanted nor tried to do.

Perhaps the fact that the sniper neither tried nor wanted to alert his enemies to his whereabouts explains why most of the participants in C1 judged that he did not intend to do so. Nevertheless, since the sniper clearly intended to fire his gun and clearly intended to kill the enemy commander—knowing all the while that he would thereby alert his enemies to his presence—participants judged that he *intentionally* alerted them. Since alerting the enemy was one of the foreseen results of carrying out the sniper's general intention to shoot the enemy commander, perhaps participants judged that the sniper intentionally alerted the enemies even though he neither wanted nor intended to do so. The same goes for the results of C2.

In both C1 and C2, the sniper's overall intention to shoot the enemy commander was apparently sufficiently related to the foreseen side effects of his actions that most participants judged that the sniper brought about these side effects intentionally. Presumably, participants judged that the sniper *x-ed* intentionally because they judged that he acted with an *x*-relevant intention—an intention robust enough to stand in for a direct intention to *x*. At the same time, because very few participants judged that the sniper wanted to bring about the respective side effects, very few of them judged that he intended to bring these side effects about.

People also appear to have a higher threshold for ascriptions of intentional action when the actions or side effects in question are morally positive than they do when the actions or side effects in question are either non-moral or morally negative. That such an asymmetry exists can no longer be doubted—even if more work admittedly needs to be done to figure out precisely what explains the asymmetry.¹⁴ It is also unclear whether the asymmetry is warranted or whether it is merely the result of misguided moral heuristics and cognitive biases (see, e.g. Nadelhoffer, 2004 c, forthcoming a; Alicke, 2000; Greene and Haidt, 2002; Haidt, 2001; Knobe and Mendlow, forthcoming; Mele, 2001; Mele and Sverdlik, 1996).

¹⁴ I have elsewhere argued that it would not be surprising if an evolutionary model along the following lines explained the asymmetry (Nadelhoffer, 2004c). Imagine that long ago humans who were good at quickly detecting 'harmers'—i.e. morally blameworthy individuals such as cheats, liars, thieves, rapists, murderers, and other scoundrels—were more apt to survive than those who did not. For humans living under these conditions, the best survival strategy would be to blame *first* and worry about exculpating or mitigating circumstances *later*. Given that this kind of survival strategy is at least possible—if not likely—it is possible that our judgments about the blameworthiness of an action may come before our determination of whether the action was performed intentionally. More importantly, this would not undermine the usefulness of the concept of intentional action to the extent that this concept—along with a cluster of other closely related concepts such as purposely, knowingly, and accidentally—could still be used to amplify, verify, mitigate, or exculpate our antecedent attributions of moral responsibility. And even though in these situations our notion of intentional action would admittedly not play its usual role of fixing blame, it would nevertheless have an important role to play—*viz.*, helping to ensure that our initial 'harmer detection' reaction was not unjustified.

These unsettled issues notwithstanding, we presently have enough evidence to conclude that while the SV may adequately capture folk intuitions concerning cases involving morally positive actions and side effects, it nevertheless fails to cohere with folk intuitions in cases involving morally neutral or bad actions or side effects.¹⁵ As such, we have reason to reject SV as a general analysis of the folk concept of intentional action. However, it remains an open question whether SPV succeeds where SV fails. On the surface at least, it appears that in each of the studies that have been run so far, the agents in the vignettes intended to do *something*, even if they happened *not* to have intended to bring about the particular outcome in question. This suggests that SPV agrees with a much broader range of folk intuitions about intentional actions than SV—which I have argued is less intuitive to laypersons than its proponents have traditionally assumed.

*Department of Philosophy
Dickinson College*

References

- Adams, F. 1986: Intention and intentional action: the simple view. *Mind & Language*, 1, 281–301.
- Adams, F. and Steadman, A. 2004a: Intentional action in ordinary language: core concept or pragmatic understanding? *Analysis*, 64, 173–181.
- Adams, F. and Steadman, A. 2004b: Intentional actions and moral considerations: still pragmatic. *Analysis*, 64, 264–267.
- Alicke, M. D. 2000: Culpable control and the psychology of blame. *Psychological Bulletin*, 126:4, 556–574.
- Bratman, M. 1984: Two faces of intention. *Philosophical Review*, 93, 375–405.
- Bratman, M. 1987: *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- Doris, J. and Stich, S. 2003: As a matter of fact: Empirical perspectives on ethics. In F. Jackson and M. Smith (eds), *The Oxford Handbook of Contemporary Analytic Philosophy*. Oxford: Oxford University Press.
- Goldman, A. and Pust, J. 1998: Philosophical theory and intuitional evidence. In M. R. DePaul and W. Ramsey (eds), *Rethinking Intuition: The Psychology of Intuition*

¹⁵ The existing data suggest that people normally judge that trying and wanting to *x* are necessary for intending to *x* even if they do not believe that trying and wanting to *x* are necessary for intentionally *x-ing*. Moreover, it appears that if an agent both wanted and intended to *x* and intentionally *x-ed* knowing (and not caring) that doing *x* would bring about some morally neutral or bad side effect *y*, then people will judge that the agent intentionally brought *y* about regardless of whether she wanted or intended to *y*. If, on the other hand, *y* is a morally *positive* side effect, people will not ordinarily judge that the agent intentionally brought *y* about unless she both wanted and intended to do so.

- and *Its Role in Philosophical Inquiry*. Lanham: Rowman and Littlefield Publishers, pp. 179–197.
- Greene, J. and Haidt, J., 2002: How (and where) does moral judgment work? *Trends in Cognitive Science*, 6, 517–523.
- Harman, G. 1976: Practical reasoning. *Review of Metaphysics*, 79, 431–463. Reprinted in A. Mele ed. *The Philosophy of Action*. Oxford: Oxford University Press, 1997.
- Haidt, J. 2001: The emotional dog and its rational tail: a social intuitionist approach to moral judgment. *Psychological Review*, 108, 814–834.
- Heider, F. 1958: *The Psychology of Interpersonal Relations*. New York: Wiley.
- Jones, E. E. and Davis, K. E. 1965: From acts to dispositions: the attribution process in person perception. In L. Berkowitz (ed.), *Advances in Experimental Social Psychology*. Hillsdale, NJ: Erlbaum, pp. 371–388.
- Knobe, J. 2003a: Intentional action and side effects in ordinary language. *Analysis*, 63, 190–94.
- Knobe, J. 2003b: Intentional action in folk psychology: an experimental investigation. *Philosophical Psychology*, 16:2, 309–24.
- Knobe, J. 2004: Intention, intentional action, and moral considerations. *Analysis*, 64, 181–187.
- Knobe, J. and Burra, A. Forthcoming: What is the relation between intention and intentional action? *The Journal of Cognition and Culture*.
- Knobe, J. and Mendlow, G. Forthcoming: The good, the bad, and the blameworthy: understanding the role of evaluative reasoning in folk psychology. *The Journal of Theoretical and Philosophical Psychology*.
- Lynch, M. 2005: Trusting intuition. In P. Greenough and M. Lynch (eds), *Truth and Realism: Current Debates*. New York: Oxford University Press.
- Malle, B. and Knobe, J. 1997: The folk concept of intentional action. *Journal of Experimental Social Psychology*, 33, 101–121.
- Malle, B. and Knobe, J. 2001: The distinction between desire and intention: a folk-conceptual analysis. In B. F. Malle, L. J. Moses and D. A. Baldwin (eds), *Intentions and Intentionality*. Cambridge, MA: MIT Press.
- McCann, H. 1986: Rationality and the range of intention. *Midwest Studies in Philosophy*, 10, 191–211.
- McCann, H. 1991: Settled objectives and rational constraints. *American Philosophical Quarterly*, 28, 24–36.
- McCann, H. 1998: *The Works of Agency: On Human Action, Will, and Freedom*. Cornell, NY: Cornell University Press.
- McCann, H. Forthcoming: Intentional action and intending: Recent empirical studies. *Philosophical Psychology*.
- Mele, A. 1992: *Springs of Action: Understanding Intentional Behavior*. New York: Oxford University Press.
- Mele, A. 2001: Acting intentionally: probing folk notions. In B. F. Malle, L. J. Moses, and D. A. Baldwin (eds) *Intentions and Intentionality*. Cambridge, MA: MIT Press.
- Mele, A. 2005: Decisions, intentions, and free will. *Midwest Studies in Philosophy*, 24, 146–162.

- Mele, A. and S. Sverdlík. 1996: Intention, intentional action, and moral responsibility. *Philosophical Studies*, 82, 265–287.
- Nadelhoffer, T. 2004a: The Butler problem revisited. *Analysis*, 64, 277–284.
- Nadelhoffer, T. 2004b: Praise, side effects, and intentional action. *The Journal of Theoretical and Philosophical Psychology*, 24, 196–213.
- Nadelhoffer, T. 2004c: Blame, badness, and intentional action: a reply to Knobe and Mendlow. *The Journal of Theoretical and Philosophical Psychology*, 24, 259–269.
- Nadelhoffer, T. 2005: Skill, luck, and intentional action. *Philosophical Psychology*, 18:3, 343–354.
- Nadelhoffer, T. Forthcoming a: Bad agents, blameworthy agents, and intentional actions: some problems for juror impartiality. *Philosophical Explorations*.
- Nadelhoffer, T. Forthcoming b: Foresight, moral considerations, and intentional actions. *The Journal of Cognition and Culture*.
- Nahmias, E., Morris, S., Nadelhoffer, T. and Turner, J. 2005: Surveying free will. *Philosophical Psychology*, 18:5, 561–584.
- Nahmias, E., Morris, S., Nadelhoffer, T. and Turner, J. Forthcoming: Is incompatibilism intuitive? *Philosophy and Phenomenological Research*.
- Nichols, S. 2004a: After objectivity. *Philosophical Psychology*, 17, 5–28.
- Nichols, S. 2004b: The folk psychology of free will: fits and starts. *Mind & Language*, 19, 473–502.
- Nichols, S. 2004c: Folk concepts and intuitions: from philosophy to cognitive science. *Trends in Cognitive Science*.
- Nichols, S., Weinberg, J. and Stich, S. 2003: Metaskepticism: meditations in ethno-epistemology. In S. Luper (ed.), *The Sceptics*. Burlington, VT: Ashgate.
- Ossorio, P. G. and K. E. Davis. 1968: The self, intentionality, and reactions to evaluations of the self. In C. Gordon and K. J. Gergen (eds), *The Self in Social Interaction*. New York: Wiley.
- Shaver, K. G. 1985: *The Attribution of Blame*. New York: Springer-Verlag.
- Sosa, E. 2005: A defense of the use of intuitions in philosophy. In M. Bishop and D. Murphy (eds), *Stich and His Critics*. Malden MA: Blackwell Publishers.
- Weatherston, B. 2003: What good are counterexamples? *Philosophical Studies*, 115, 1–31.
- Weinberg, J., Nichols, S. and Stich, S., 2001: Normativity and epistemic intuitions. *Philosophical Topics*, 29, 429–460.
- Williamson, T. 2004: Philosophical ‘intuitions’ and skepticism about judgment. *dialectica*, 58, 109–153.