

Intentional Action and the Praise-Blame Asymmetry

Abstract. Recent empirical research conducted by Joshua Knobe has uncovered two asymmetries in our judgments about intentional action and moral responsibility. First, people are more inclined to say that a side effect was brought about intentionally when they regard that side effect as bad than when they regard it as good. Second, people are more inclined to ascribe blame to someone for bad effects than they are inclined to ascribe praise for good effects. These findings suggest that the notion of intentional action has a normative component. In this paper I propose a theory of intentional action on which one acts intentionally if one fails to be motivated to avoid a bad effect. This explains the asymmetry concerning intentional action. The praise-blame asymmetry is explained in terms of the claim that praise depends on being appropriately motivated, whereas blame does not.

Frank Hindriks (f.a.hindriks@rug.nl)
University of Groningen

Published in 2008, *Philosophical Quarterly* 58, 233: 630-41.

Intentional Action and the Praise-Blame Asymmetry

Joshua Knobe (2003a, 2003b, 2004a) has discovered an asymmetry in our judgments about praise and blame in relation to the side effects that our actions generate. When a side effect of someone's action is harmful, we are far more inclined to blame that person than we are inclined to praise someone when the side effect is beneficial – even though in neither of the two cases she cared about the occurrence of the effect. I believe there is a simple explanation for this praise-blame asymmetry, which many regard as puzzling. When the side effect has positive moral worth, the person does not deserve any praise, because – by hypothesis – the fact that her action would have that effect did not play any role in her reasoning. After all, she did not care about the occurrence of the effect. When the side effect has negative moral worth, she does deserve blame, because the fact that her action would have the harmful effect should have played a role in her reasoning while it did not.

Knobe also reports an asymmetry in our judgments concerning intentional action. We are more inclined to judge that an effect is brought about intentionally when the effect is harmful as compared to when the effect is beneficial. This second finding can be explained in a way that resembles the proposed explanation of the praise-blame asymmetry, or so I shall argue. Concerning the beneficial effect we judge that it is not brought about intentionally because the effect did not play any role in the practical reasoning of the person at issue. In relation to the harmful effect, we do conclude that it is brought about intentionally. In this case, the person ignores a consideration she should treat as a reason against performing the action. Again, the fact that the relevant person does not care about bringing about the effect is crucial to the explanation of this asymmetry concerning our judgments about intentional action.

My hypothesis, then, is that both asymmetries can be explained in terms of the considerations that an agent takes or should take into account when thinking about what to do, i.e. in terms of motivating and normative reasons (section 1). The key implication of the explanation offered here is that the notion of intentional action has a normative component: it should partly be explicated in terms of normative reasons. I present a theory of intentional action that accommodates this idea in a natural way (section 2). I end the paper by contrasting my view with explanations others have offered (section 3). Some have argued that the asymmetry in our judgments about intentional action is to be explained in terms of the asymmetry in our judgments about praise and blame, this in spite of the fact that common

sense has it that judgments about intentional action are in some sense prior to those about moral responsibility. My explanation is more general in that it can explain why both views can seem right, even though both are wrong. Knobe claims that when making judgments about intentional action we attend to different features ‘depending on whether the behavior itself is good or bad’ (2006, 225-26). But he does not tell us which features they are. As the explanation offered here facilitates the identification of the relevant features, it has more depth than Knobe’s explanation.

1. Explaining the Asymmetries

1.1 The Asymmetries

Consider what I shall call the chairman experiment. Subjects were randomly assigned to a version of the chairman scenario that includes either a help condition or a harm condition. Subjects in 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, 191)

Subjects in the help condition read the same vignette except for the fact that ‘harm’ was replaced by ‘help’. Subsequently, they were asked how much blame (praise) the chairman deserved for the effect his decision had on the environment (on a scale from 0 to 6). They also answered the question whether they thought the chairman intentionally harmed (helped) the environment.

Most of the subjects who read the vignette with the harm condition said that the chairman brought about the side effect intentionally (82%), whereas most of those who were confronted with the help condition said that the chairman did not bring about the side effect intentionally (77%). These results are statistically significant, so they reveal a stark asymmetry between our judgments about intentional action in relation to beneficial versus

harmful effects. I shall call this ‘the Intentional Action Asymmetry’. In addition to this, subjects said that the chairman deserved a lot of blame in the harm condition, but very little praise in the help condition (these results are statistically significant as well; the mean values are 4.8 and 1.4 respectively). I shall call this ‘the Praise-Blame Asymmetry’. Finally, the total amount of praise or blame that subjects offered was significantly correlated with their judgments about whether or not the side effect was brought about intentionally.

1.2 Explaining the Praise-Blame Asymmetry

In order to explain the Praise-Blame Asymmetry, we should note that neither in the help nor in the harm version of the scenario the chairman cares about the environment. This fact stops us from praising him in the help scenario, while it does not stop us from blaming him in the harm scenario. The underlying idea is, roughly, that praise requires effort on the part of the person who receives it, while blame does not. We only praise someone for bringing about a morally good consequence if she performed the relevant action in order to bring about this effect. If that is the case, the effort expended in performing the action was made at least in part for this purpose. There is no such requirement for blame. A failure to pay attention to a morally bad consequence of one’s action is perfectly compatible with blame. The Praise-Blame Asymmetry, then, can be explained in terms of the following contrast: people need to be motivated appropriately in order to receive praise, but not for blame. All there is to the praise-blame asymmetry is that we do not praise people for bringing about a beneficial effect when they do not care about it, whereas we do blame people for bringing about a harmful effect when they do not care about it.¹

The point can also be made in terms of motivating and normative reasons. Motivating reasons play a role in explaining actions, while normative reasons serve a justificatory role. More precisely, a motivating reason is a consideration one actually takes into account in that one treats it as reflecting favorably on the action one ends up performing. A normative reason

¹ McCann provides a similar explanation of why the chairman is not praised in the help version of the scenario: ‘And once the chairman disavows this interest, we have no reason to consider him praiseworthy for helping the environment.’ (2005, 745) However, he maintains that the chairman intends to harm the environment in the harm version, in spite of the counterevidence he presents (ibid., 740-41). On my view he merely does so intentionally. (See also notes 4 and 9.)

is a consideration that counts in favor or against an action and that one should take into account when considering it, irrespective of the question whether one does.² Normative reasons often pertain to the effects of our actions. This can be made explicit in terms of what I call the Side-Effect Deliberation Norm, SDN, which appears to govern our deliberation:

[SDN] When a side effect is morally significant, one should take it into consideration while deliberating about what to do.

Violation of SDN blocks praise, but does not block blame. When the side effect is morally bad, a violation of SDN justifies blame. After all, one fails to take into account a consideration one should pay attention to, i.e. a normative reason. When the side effect is morally good, a violation of SDN stops us from praising the agent. Praise is justified only when one is motivated by the considerations that one should take into account, i.e. only when one's motivating reasons are suitably aligned with the normative reasons that bear on the action at issue.

As he does not care about the environment, the chairman violates SDN in both versions of the scenario. So, in neither version of the scenario does he assign weight to the consequence of his action, even though it is morally significant. At the same time, we believe he should have attached significance to the effect the business strategy has on the environment. And we assume that he is aware of the normative significance of the environment. Maybe businessmen have a duty not to harm the environment, but no duty to help it, at least not if this goes at the expense of the amount of profit made. If so, opportunities to help the environment only need to be taken up when they have no (or hardly any) effect on the amount of profit a business will make. This means that in many situations in which a business strategy is expected to benefit the environment businessmen hardly need to stop and think about it. Not caring, however, is never the appropriate attitude. So, his lack of

² Motivating reasons are often taken to be psychological states (Smith 1994). I do not mean to rule out the possibility that they are indeed psychological states by using the term 'consideration' for such reasons. The characterization in the main text is merely meant to serve as a means for getting across the basic idea underlying the distinction between normative and motivating reasons. Similar distinctions are made using different terms including 'justifying' versus 'motivating' (Frankena 1958) and 'normative' versus 'operative' (Scanlon 1998).

appropriate motivation explains the lack of praise. Our blame is explained by the fact that he acts while ignoring an important moral consideration that he should treat as a reason against implementing the strategy. The puzzlement that contributors to the debate – including Adams and Steadman, Knobe, McCann, and Nadelhoffer – have expressed in response to the praise-blame asymmetry suggest that they expected the chairman to be praised for helping the environment as much as he is blamed for harming it. His violation of SDN explains why this is not what we see. Since he does not care about the environment, the chairman’s reasoning is morally flawed both in the harm and in the help versions of the scenario (in this respect they are symmetrical!).

1.3 Explaining the Intentional Action Asymmetry

The distinction between motivating reasons and normative reasons also provides the key to the explanation of the Intentional Action Asymmetry. It is generally agreed that, when one performs an action (in part) in order to bring about a certain effect and one’s action does indeed generate the effect in the expected manner, the effect is brought about intentionally. This means that an effect can be brought about intentionally due to the fact that it functions as a motivating reason for performing the action that generates the effect, or due to the fact that one ascribes positive significance to the expected occurrence of the effect. What I propose to add to this in order to explain the asymmetry is the idea that one brings about an effect intentionally if one performs the intended action in spite of the fact that one believes that one ought to treat the effect as a reason not to perform the intended action. When this is the case, one should ascribe negative significance to it, but one fails to do so. This implies that one can bring about an effect intentionally even if one ascribes no significance to it in one’s deliberation. In other words, an effect can be brought about intentionally due to the fact that its expected occurrence provides a normative reason against performing an action and one ignores this reason.³

How can these ideas be used for explaining the Intentional-Action Asymmetry? Again we should focus on the fact that the chairman does not care about the effect he brings about.

³ This appears to be one of the things that Harman is after when he claims that ‘[o]ne can do something intentionally even though one does not intend to do it, if one does it in the face of what ought to be a reason not to do it and, either one tries to do it, or one does it as a foreseen consequence of something else that one intends to do’ (1976, 434).

In the help version of the scenario, the fact that the chairman is not motivated to help the environment explains why we say he does not do so intentionally. In the harm version, the chairman fails to be moved by a consideration to which he should ascribe negative significance, the harm that will be done to the environment when he pursues the profit-maximizing strategy. The suggestion I just made entails that, because of this, the chairman harms the environment intentionally. After all, he ignores a normative reason against performing the action. So, the Intentional-Action Asymmetry is explained by the fact that in the help version of the scenario the chairman fails to have the motivating reason which would have made his act of helping the environment an intentional one on the one hand, and the fact that in the harm version of the scenario the chairman fails to be motivated in the way he should be.

The explanation of our judgments in the harm version can be made more precise in terms of what I shall call ‘the Deontic Significance Thesis’ [DST]:

An agent *S* *As* intentionally if *S* intends to *B*, *As* by *B*-ing, expects to *A* by *B*-ing, and *S* *Bs* in spite of the fact that she believes her expected *A*-ing constitutes a normative reason against her *B*-ing.

This means that ignoring a normative reason or failing to be motivated to avoid a wrongdoing lies at the heart of a sufficient condition for acting intentionally. In the next section I shall argue that DST is not ad hoc but fits seamlessly into an overall account of intentional action.

2. A Theory of Intentional Action

According to what Bratman (1984, 1987) has called the Simple View of intentional action, one intentionally *As* only if one intends to *A*. Knobe’s experiments reveal that this view is implausible. The chairman does not intend the harm his action causes (see McCann 2005, 740-41 for data that support this).⁴ The main alternative to the Simple View is the Single Phenomenon View, according to which intending and acting intentionally involve a common

⁴ In spite of their apparent message, McCann tries to use these data to bolster the case for the Simple View (see also Adams and Steadman 2004a, 2004b). See Knobe (2006) and Nadelhoffer (2006) for critiques.

state. I shall follow Bratman in assuming that this state is that of an intention. The idea is that one intentionally *As* only if one intends to *B*, where *B* may (but need not) be distinct from *A*. This view can be developed in several ways. In this section, I shall formulate a version of the Single Phenomenon View that accounts for the Intentional Action Asymmetry.

One way in which one can *A* intentionally is, of course, by intending to *A*. *A*-ing may be the only thing one aims at. However, one can also intend to *A* if one expects to *A* in the course of performing another action *B* that one intends. This will be the case if one *Bs* in part because one expects to *A* in the course of *B*-ing. In other words, one ascribes positive significance to the fact that one's *B*-ing will lead to one's *A*-ing while deliberating about whether to *B*. In such cases, one's expectation to *A* in the course of *B*-ing is a motivating reason for one's *B*-ing. Suppose, for instance, that the chairman does care about the effect that the business strategy has on the environment. In relation to the help condition, this means he supports the profit-maximizing strategy in part because it helps the environment. This allows us to say he supports it in order to benefit the environment. And this implies that he intends to benefit the environment, which means of course that he does so intentionally. In such cases acting intentionally is a matter of ascribing positive significance to an effect of the action at issue. Note that, when one ascribes positive significance to its expected occurrence, an effect ceases to be a side effect and becomes an intended effect, and the act of bringing it about becomes an intended action. I shall call the claim that one *As* intentionally if one *Bs* and one intends to do so in part because one expects to *A* by *B*-ing the Positive Significance Thesis [PST].

PST reveals that intentional action is intimately connected with reasons for action. On the theory of intentional action presented and defended here, this is what unites all instances of acting intentionally: intentional action concerns the reasons we consider or fail to consider while deliberating about what to do. PST as such is consistent with the Simple View of intentional action. In order to account for the Intentional Action Asymmetry, however, this view has to be abandoned. Those who reject the Simple View thereby embrace the idea that one can bring about an effect intentionally without intending to bring it about – an idea that received its first systematic defense in Anscombe (1957). If one intends to bring about an effect, one has a reason for doing so, which may just be that one wants to do so. I might, for instance, board a plane in order to get to New York. The state of affairs of me being in New York is the intended effect of my action. It appears that bringing about an effect intentionally is not only consistent with having a reason that counts in favor of doing so, but also with having a reason against doing so. The purpose of my trip to New York might be to visit a

friend. Now, my friend might live in a dangerous part of New York. As I do not want to go to dodgy places, the fact that she lives in a dangerous area is a reason for me not to visit her. However, as it happens, my wanting to see her wins out. It seems that in this case I go to the dangerous neighborhood intentionally, because I visit her in spite of the fact that she lives in a dangerous neighborhood. As I have no reason for doing so, I do not intend to do so. The only thing I intend to do is visiting my friend.

Gilbert Harman's well-known sniper example supports the idea that such a condition is indeed relevant to doing something intentionally. Harman writes: 'In firing his gun, the sniper knowingly alerts the enemy to his presence. He does this intentionally, thinking that the gain is worth the possible cost. But he certainly does not intend to alert the enemy to his presence.' (1976, 433; emphasis added) The unintended effect of the sniper's firing his gun is that the enemy is alerted to his presence. His action of alerting the enemy is performed intentionally because he expects that he will alert the enemy by firing his gun, and he fires his gun in spite of the fact that he does not want to alert the enemy. In cases like this, the agent intends to *B*, does not intend to *A*, As by *B*-ing, and expects to do so. In addition to this, the agent intends to *B* in spite of the fact that she does not want to *A*, or – what is meant to be equivalent – that she takes the fact that she will *A* by *B*-ing to be a reason against her *B*-ing. I call the claim that these conditions suffice for *A*-ing intentionally the Negative Significance Thesis [NST].⁵

In the previous section, I already suggested that one might do something intentionally without intending it when one does it in spite of the fact that one has a normative reason against doing it. Suppose that, when in New York, I decide to take a taxi. All I care about is arriving at my friend's house at the time we agreed upon, and I am running late. I see a pregnant woman trying to get to the taxi rank. Apparently she is in early labor. As I can move quicker, I get to the taxi rank first, and take the only taxi available at that time. Presumably, I ought to have given way to the pregnant woman and help her to get into the taxi. Instead, I prevent her from taking the taxi to get to the hospital. As I know I should care about the woman, I do so intentionally. After all, I ignore a normative reason that counts against taking the taxi myself. Cases like this are similar to NST-cases in that they concern situations in which the agent intends to *B*, does not intend to *A*, As by *B*-ing, and expects to do so. What is distinctive about them is that the agent intends to *B* in spite of the fact that she believes she

⁵ See Bratman (1987, 123) and Knobe (2006, 220) for other examples that are accommodated by the NST.

ought to treat the fact that she will *A* by *B*-ing as a reason against her *B*-ing. The Deontic Significance Thesis [DST] is the claim that this suffices for *A*-ing intentionally. It serves to account for cases in which an agent does not ascribe any significance to an effect but still brings it about intentionally.

The upshot is that there are at least three ways in which one can do something intentionally captured by the three theses presented above – the PST, the NST, and the DST. First, by wanting to do the action that one expects to do in the course of performing an intended action. Second, by performing the intended action in spite of not wanting to do the action that one expects to do in the course of performing it. Third, by performing the intended action in spite of the fact that (one believes that) the unintended action provides a normative reason (which need not be overriding) for not performing the intended action. These can be combined into an analysis of intentional action, [IA]:⁶

- [IA] An agent *S* *As* intentionally if *S* intends to *B*, *As* by *B*-ing, expects to *A* by *B*-ing, and
- (a) *S* intends to *B*, because she expects thereby to *A*, or
 - (b) *S* *Bs* in spite of the fact that she does not want to *A*, or
 - (c) *S* *Bs* in spite of the fact that she believes her expected *A*-ing constitutes a normative reason against her *B*-ing.

According to [IA] one's *A*-ing intentionally can be due to one's ascribing positive or negative significance to it when deliberating about *B*-ing, or to the fact that one should ascribe negative significance to it. This analysis makes it easy to see why our judgments concerning intentional action are asymmetric when the cases considered involve morally significant

⁶ This analysis of intentional action is incomplete in at least three respects. First, it does not cover cases in which one's *A*-ing is a reason for *B*-ing without this implying that one intends to *B* (see, for instance, Bratman's well-known video-games example; 1987, 113-15). Second, it does not accommodate the way in which skill and deviant causal chains influence our judgments concerning intentional action. Third, it does not account for experimental results concerning luck (see, for instance, Nadelhoffer 2004a). I plan to address cases involving luck in future work. See Mele and Moser (1994) for an excellent discussion of how skill and causal chains should enter the analysis. Bratman (1987, 121) formulates a condition that covers cases like his video-games example that is compatible with [IA].

effects about which the relevant agent does not care. Because of the fact that one ought to assign weight to harmful effects in one's reasoning, condition (c) can be satisfied even though the agent at issue does not care about the effect. However, even though one ought to assign weight to beneficial effects in one's reasoning as well, there is no condition analogous to (c) that can apply when one does not care about a beneficial effect. Instead, failure to satisfy condition (a) explains why we say one does not do it intentionally.

Accepting DST entails embracing the idea that intentional action should not be analyzed in terms of motivating reasons only, but in terms of normative reasons as well. In other words, intentional action has a normative component. Even though the idea that our concept of intentional action includes a normative condition is very controversial (Mele and Sverdlik 1996), it is quite intuitive on the line I defend here. After all, this way of acting intentionally is a matter of ignoring a consideration that should be treated as a reason against implementing the strategy, or of knowingly violating SDN. The fit with the observed asymmetries provides further support for this idea.

3. Intentional Action and Moral Responsibility

We commonly suppose that whether or not someone performed an action intentionally is relevant to moral responsibility. For instance, having performed a harmful action accidentally rather than intentionally can function as a mitigating circumstance. When it does, one should receive less blame as compared to a similar situation in which one performed the action intentionally. Most contributors to the literature on the Praise-Blame Asymmetry and the Intentional Action Asymmetry seem to hold or have held the view that common sense has things backwards here. Our judgments concerning moral responsibility feed into our judgments about intentional action rather than the other way around.⁷ Malle and Nelson (2003) argue that judgments about blame bias our judgments about intentionality. Nadelhoffer (2004b) maintains that our judgments about praise and blame can influence our judgments

⁷ One can easily get the impression that Knobe also believes that the Praise-Blame Asymmetry explains the Intentional Action Asymmetry. Consider, for instance, the following passage: 'And this asymmetry in people's assignment of praise and blame may be at the root of the corresponding asymmetry in people's application of the concept intentional' (Knobe 2003a, 193; emphasis in original). However, he explicitly argues against this idea in Knobe (2004b) and Knobe and Mendlow (2004).

about intentional action, but also allows for the moral character of the effects to do so. Adams and Steadman (2004a and 2004b), as well as McCann (2005) defend the view that denying that the chairman harms the environment intentionally carries with it the pragmatic implicature that he should not be blamed. As most of the people questioned want to blame the chairman, they say instead that he did harm the environment intentionally. Some have gone as far as claiming that the Praise-Blame Asymmetry in combination with the fallacious belief that blame requires intentional action explains the Intentional Action Asymmetry (Mele 2001).⁸ An important argument in favor of the explanation I offer in this paper is that it can account for both of these perspectives on the relation between intentional action and moral responsibility.⁹

In effect, I have argued that both asymmetries have a common denominator: motivating and normative reasons for action (see sections 1.2 and 1.3). And when two factors depend in an analogous way on a third one, it is often easy to be misled in thinking that one of them is in some sense prior to the other, or vice versa. Suppose, for instance, that poor people watch more television including violent programs than others, and that they also commit more violent crimes. This information by itself is insufficient for determining whether watching violent programs causes or increases the incidence of violent crimes. Instead, it might be that both can be traced back to poverty. The claim that is implicit in the analysis defended here is that there is no order of priority between intentional action on the one hand and moral responsibility on the other. Instead, what matters to both is what motivated the agent and what should have motivated her.

⁸ Knobe (2006, 213-14) discusses evidence that conflicts with this view. Mele has since then retracted it (see his 2003, 333).

⁹ As we saw in note 4, Adams and Steadman as well as McCann try to save the Simple View according to which one performs an action intentionally only if one intends that action. The strongest argument against this view is that people often say that an action is performed intentionally while also claiming that it was not intended (McCann 2005, Nadelhoffer 2006). The point of the comparison offered here, then, is not to provide further arguments against the Simple View. Instead, it is to bolster my claim that [IA] provides a plausible explanation of the Intentional Action Asymmetry. In addition to being a natural extension of existing theories of intentional action, it can be used to make sense of the ways in which people have responded to the asymmetries.

How does my explanation of the asymmetries compare to that provided by Knobe? As indicated earlier, Knobe holds that the moral character of the effect can influence our judgments about intentional action (which is, of course, also relevant to an adequate explanation of the Praise-Blame Asymmetry). This is born out by my explanation. Consider DST once more. It refers to a normative reason not to perform a particular action or bring about a particular effect. This implies that bringing about the effect as such is bad (although there might be circumstances in which one should do so). In addition to this, Knobe (2006) has suggested that depending on the moral character of the behavior, different features might be relevant for determining whether or not an action was performed intentionally. Although he has mentioned several candidate features, Knobe has not provided anything like a systematic account that explains the Intentional Action Asymmetry. The theory of intentional action presented in the previous section does just this. Hence, it provides for a deeper explanation.

[IA] facilitates the identification of the features that explain the asymmetry. In all cases, the effect is expected or foreseen. When the side effect is good, people wonder whether the agent treats this as a reason for doing *B*. The question is whether the agent is appropriately motivated (which presumably entails that she tries to bring about the effect). In the experiments it is assumed that this is not the case (recall that the chairman is assumed not to care about the environment¹⁰). So, people judge that the agent does not *A* intentionally. When the side effect is bad, people focus on the fact that the agent ought to treat it as a reason against *B*-ing (given DST the relevant factor is whether he is aware of this, but it appears that people presuppose that he is when they form their judgments). After all, whether the agent treats it in this way is irrelevant to the question of whether he does it intentionally. So, when the behavior is good, we attend to the considerations the agent actually takes into account, while when the behavior is bad we focus on what he should consider, irrespective of whether he does. The overall conclusion to be drawn from the two asymmetries, then, is that the notion of intentional action should not only be analyzed in terms of motivating reasons, but also in terms of normative reasons.¹¹

¹⁰ An analogous assumption is made in the lieutenant experiment reported in Knobe (2003a).

¹¹ I am grateful for valuable comments from Michael Bratman, Martin van Hees, Joshua Knobe, Arto Laitinen, Teemu Toppinen, and two anonymous referees.

References

- F. Adams and A. Steadman, 'Intentional Action in Ordinary Language: Core Concept or Pragmatic Understanding', Analysis, 64 (2004a), pp.173-81.
- F. Adams and A. Steadman, 'Intentional Actions and Moral Considerations: Still Pragmatic', Analysis, (2004b) 64, pp.264-67.
- G.E.M. Anscombe, Intention (Oxford: Basil Blackwell, 1957).
- M. Bratman, 'Two Faces of Intention', Philosophical Review, 93 (1984), pp.375-405.
- M. Bratman, Intention, Plans, and Practical Reason (Cambridge (MA): Harvard University Press, 1987).
- W. Frankena, 'Obligation and Motivation in Recent Moral Philosophy', in A.I. Melden (ed), Essays on Moral Philosophy (University of Washington Press, 1958), 40-81.
- G. Harman, 'Practical Reasoning', Review of Metaphysics, 29 (1979), pp.431-63.
- J. Knobe, 'Intentional Action and Side Effects in Ordinary Language', Analysis, 63 (2003a), pp.190-94.
- J. Knobe, 'Intentional Action in Folk Psychology: An Experimental Investigation', Philosophical Psychology, 16 (2003b), pp.309-24.
- J. Knobe, 'Intention, Intentional Action and Moral Considerations', Analysis, 64 (2004a), pp.181-87.
- J. Knobe, 'Folk Psychology and Folk Morality: Response to Critics', Journal of Theoretical and Philosophical Psychology, 24 (2004b), pp.270-79.
- J. Knobe, 'The Concept of Intentional Action: A Case Study in the Uses of Folk Psychology', Philosophical Studies, 130 (2006), pp.203-31.
- J. Knobe and G. Mendlow, 'The Good, the Bad, and the Blameworthy: Understanding the Role of Evaluative Reasoning in Folk Psychology', Journal of Theoretical and Philosophical Psychology, 24 (2004), pp.252-58.
- B.F. Malle and S.E. Nelson 'Judging Mens Rea: The Tension between Folk Concepts and Legal Concepts of Intentionality', Behavioral Sciences and the Law, 21 (2003), pp.563-80.
- H.J. McCann, 'Intentional Action and Intending: Recent Empirical Studies', Philosophical Psychology, 18 (2005), pp.737-48.
- A. Mele, 'Acting Intentionally: Probing Folk Notions', in B. Malle, L. Moses, and D. Baldwin (eds), Intentions and Intentionality: Foundations of Social Cognition (Cambridge (MA): MIT Press, 2001), 27-43.

- A. Mele, 'Intentional Action: Controversies, Data, and Core Hypotheses', Philosophical Psychology, 16 (2003), pp.325-40.
- A. Mele, and P. Moser, 'Intentional Actions', Nous, 28 (1994), pp.39-68.
- A. Mele and S. Sverdlik, 'Intention, Intentional Action, and Moral Responsibility', Philosophical Studies, 82 (1996), pp.265-87.
- T. Nadelhoffer, 'The Butler Problem Revisited', Analysis, 64 (2004a), pp.277-84.
- T. Nadelhoffer, 'Blame, Badness, and Intentional Action: A Reply to Knobe and Mendlow', Journal of Theoretical and Philosophical Psychology, 24 (2004b), pp.259-69.
- T. Nadelhoffer, 'On Trying to Save the Simple View', Mind & Language, 21 (2006), pp.565-86.
- T.M. Scanlon, What We Owe to Each Other (Cambridge (MA): Harvard University Press, 1998).
- M. Smith, The Moral Problem (Oxford: Blackwell, 1994).