

Intentions Confer Intentionality upon Actions: A Reply to Knobe and Burra

Frederick Adams
University of Delaware

Journal of Cognition and Culture

Is intentionally doing A linked to the intention to do A? Knobe and Burra believe that the link between the English words 'intention' and 'intentional' may mislead philosophers and cognitive scientists to falsely believe that intentionally doing an action A requires one to have the intention to do A. Knobe and Burra believe that data from other languages shows convincingly that the link between intending to A and doing A intentionally is accidental and comes apart. I deny that those who believe that intentionally doing A requires the intention to do A ever thought so because of the etiology of the roots or the morphology of the words, and that intentionally doing A and intending to do A are nomologically tied for non-linguistic reasons.

Knobe and Burra (this journal) are stalking an answer to the question "What is the relation between intentions and intentional action?" My answer is that intentions cause and sustain the very actions the intentions make intentional. Actions are made intentional by the intentions to do them. Knobe and Burra challenge¹ this idea, but not successfully in my view.

Knobe and Burra also say that the concepts of intention and intentional action are important to folk psychology. And, of course, they are right. We tend to understand each other's behavior and to interact with one another on the basis of our folk psychological attributions of mental states to one another. We also hold people morally and legally responsible for their actions based upon concepts of actions and mental states that operate at the folk psychological level of precision. We hold people responsible for things they do on purpose and for harms they knew (or should have known) they would or might

cause. So the basic concepts of *intention* and *intentional action* operate at a fairly basic level of understanding in daily life. What Knobe and Burra want to challenge is that the concepts always *should* or even always *do* come as a *matched set*. They argue that the very concepts of *intention* and *intentional action* can and *do* come apart. I want to suggest why I don't believe that they can come apart, in spite of the sometimes *surprising* and sometimes *stunning* data Knobe and Burra are able to offer in support.

I: Nomological link between intending with intentional behavior

First let me say something about why intentions (a type of mental state) and actions which are intentionally performed (a kind of behavior) can not come apart. It has *nothing to do with* the fact that in English the words 'intention' and 'intentional' are related by root origin or by morphology (as suggested by Knobe and Bura). Then I'll turn my attention to the specifics of the challenge by Knobe and Burra.

Purposive behavior, whether it be the behavior of intentional agents, or of animals, insects, or even man-made artifacts (homing missiles), must be directed at an end or goal. To be purposive is the opposite of being random and just ending up at a happy place. To be purposive is to be directed at a goal by some kind of design or selection. Caterpillars crawl toward the light because light and food are correlated in normal environments. They must have a mechanism for detecting the light and its direction (which they do) and this light detection mechanism must causally guide them to the light (and the food) if their behavior is purposive (and not a lucky accident of getting to the food).

¹ They challenge it here and Knobe has a series of papers (Knobe, 2003a, 2003b, 2004) making similar challenges. I defend the idea here and elsewhere Adamd (forthcoming), Adams & Steadman (manuscript, 2004a,2004b).

Homing missiles as well can be designed to seek a heat source. The detection of a source of heat can be fed into the guidance mechanisms and as long as the missile is set to change its direction to minimize the difference between its current location and that of a heat source. It will “home in” on the heat source (even if the heat source is one of ours, not the enemies). But the point is that in neither of these systems (one natural, one artificial) are there intentions. The systems are too primitive for that. Intentions are sophisticated states of minds with rich conceptual capacities. These primitive examples of goal-directed systems don’t have that. Still, there is something in these systems, primitive as they may be, that represents or sets the end state at which the behavior of the systems is directed (light/food, in the case of the caterpillar, and heat sources in the case of the heat-seeking missile). In addition to the representations of the end-states to be reached, there must also be a feedback channel and causal dependence of output on information input and the internal guidance system, if the behavioral output of these systems is truly “purposive” and not random. Since the creation of the fields of *cybernetics* and *information theory* in the early to mid part of the 20th Century, it has been well understood how purposive systems use information and feedback mechanisms.² So the basic principles of guidance systems has been well understood for a long time and quite independently of the application of these principles to the understanding and the explanation of human behavior—which can be purposive as well, as you all know fully well. My point here is that there is, by the nature of purposive systems in general, a connection between a system having a representation of a kind of end-state to be achieved and the non-random achievement of that end state. There is, if you will, a connection between the inner workings of the mechanism and the outward behavior it

² See my (Adams, 2003)

produces or achieves. The inner mechanism causes and guides or sustains the behavior to its appropriate goal or end state. This is so in all purposive systems, no matter what words are chosen (or in what language) to represent the inner state that does the representing, the causing, and the guiding of the resultant behavior that is non-random and purposely achieved. So no part of my view has ever been that there is something special about the roots or morphologies of words used for the representing and causing and guiding of a purposive system and the word for the purposive behavior that results from this guidance system.

Second let me say how this applies to human intentional action. Human purposive behavior fits the same cybernetic model of purposive behavior as that of the caterpillar and guided missile. The difference, of course, is that we have minds and a rich set of concepts and experiences. Still for our behavior to be purposive the goals of our behavior must be represented by states of mind. Movements towards those goals must be caused by those states of mind, and the successful outcome of that behavior must causally depend in the right way on those mental representations of the goals or end-states. Otherwise, our behavior may by luck or accident achieve an end we like, but not on purpose.

Intentional behavior is behavior that we produce on purpose. It is represented in the mind and caused by those representations in response to information feedback about progress or lack of progress toward our intended goals. When I intentionally type a capitalized letter, say the letter 'S,' it is because I intend to and that intention causes my hitting the "caps" key and then the 'S' key. From time to time my left pinky hits the caps key by mistake and then I may capitalize a letter 'S,' by mistake. It is by mistake even if

I need a capitalized letter ‘S,’ if I wasn’t trying to capitalize *just then*. It is only intentional behavior if it is represented and caused by that representation in the right way (no tricks, no deviant causal chains, no accidental key presses). The mental states that represent our purposes are *intentions*. Please note that even in English there is no morphological similarity between ‘intend’ and ‘purposive,’ yet purposive behavior is caused and sustained by intentions in humans. It is at least in part because (and only because) I intend to type a capital ‘S’ that I type one intentionally here and now (or at least, in this sentence).

Those are my reasons for thinking that there is a causal and nomological link between intention and intentional action. Intentions are the mental states that represent our goals or purposes. They cause and guide our behavior to the intended ends or goals that are our purposes. When all goes well they guide and sustain our behavior to its successful outcome. So my reason for linking intentions with intentional behavior is that a cybernetic model of purposive behavior tells me that this is how purposive systems (be they natural or man-made) produce purposive behavior. This is the causal mechanism by which purposive behavior is rendered. In humans, such behavior is *intentional* and the mental states that represent the desired outcomes are *intentions*. However, it seems to me to be largely a matter of convention and arbitrary what these things are called.

II Knobe’s Original Data

Before I get to the new data based on Hindi collected by Knobe and Burra, let me rehearse some of Knobe’s findings (in English) for thinking that *intentions* and *intentional behavior* can and do come apart. I’ll discuss what I think are the limits that the data from Knobe’s original experiments show. In the next section, I’ll discuss

whether the new data of Knobe and Burra from the Hindi language shed new or different light on Knobe's original data in English. I'll argue that it does not.

Since Knobe & Burra summarize this data in this issue of the journal, I'll just remind readers of what Knobe previously found. Knobe conducted surveys in which subjects were told that a CEO of a large company knows that a new program in the company will increase profits but harm the environment. They were told that the CEO says he doesn't care one way or another about the environment, but he wants to make new profits so they should start the program. Then subjects were told that the program did increase profits and did harm the environment. When asked whether the CEO harmed the environment intentionally, 87% of subject said that he did. When asked if the CEO intended to harm the environment, only 29% said that he did. Knobe took this as support for the view that folk psychology permits judgments of actions that are *intentional*, but *not intended*. A further surprise is that when different subjects were told the same thing except that the CEO was told his new program would "help" rather than "harm" the environment (and the rest as before...the CEO doesn't care and the environment is helped), only 20% said the CEO intentionally helped the environment and no one said the CEO intended to help the environment. That is the basic pattern of replicable data Knobe found subjects typically to exhibit.

To my mind there are two stunning facts turned up by Knobe.³ One is that there is an incredible asymmetry between the cases where the subjects are judging the CEO for helping vs. harming the environment. They will blame him for harming the environment, when judging him not to intend the harm. They won't praise him for helping the

³ These are quite robust results, easily replicable. I've replicated them myself (Adams & Steadman, manuscript) as has Knobe (2003b) and McCann (forthcoming) and many others.

environment, when judging him not to intend the help. Two is that only 29% of subjects say the CEO intended harm to the environment. While a full 87% of the same subjects say that the CEO intentionally harmed the environment. In a series of papers,⁴ Annie Steadman and I have tried to explain these matters by appeal to the pragmatics of intentional talk, rather than by saying that it is true that one can do something intentionally with out intending to do it. Given what I've said above, it will be clear why I'm reluctant to say the latter.

On the pragmatics account, use of the word 'intentionally' may go beyond the literal meaning of the term. "S did A intentionally" may literally mean only that S intended to do A and that this intention caused and sustained S's doing A in the normal way, regardless of whether doing A was good or bad, right or wrong. But use of the word 'intentionally' has a further pragmatic use, as when someone steps on your foot and you say "You did that intentionally, didn't you?" Here the meaning is to scold, blame, and discourage similar future acts. Since use of intentional language has this further pragmatic dimension, I think it is this pragmatic use that best explains Knobe's data.

So to explain the dramatic asymmetry between cases where the CEO knows he will help the environment (but does not care) and where he will harm the environment (but does not care), the folk's use of "intentional" is entirely predictable on the pragmatic account. To say he helped the environment "intentionally" conversationally would be *to praise* his actions. We know that if you do some good thing *intentionally*, that is more praiseworthy than if you do it quite by accident or without caring. Yet, since the CEO professes "not to care at all" about helping the environment and says he "just wants to make a profit" the subjects don't want to credit him with a good deed. To say he did it

⁴ Adams & Steadman (manuscript, 2004a, 2004b).

intentionally would so credit him. So only 20% of subjects say the CEO helped the environment intentionally. The surprise is that there are even that many giving him credit. And no one wants to say he intended the help since he professes not to care one way or the other about the environment.

The explanation of why 87% say the CEO “intentionally” harmed the environment while only 29% say he “intended” it, again falls in line with the pragmatic account. We know that if someone does something wrong intentionally, that is worse than if one does wrong accidentally. Subjects don’t like the fact that the CEO is knowingly harming the environment and they want to blame him for it. They know that if they say he did it “intentionally” that is more blame than if he did it unintentionally. And, after all, he was told that he would harm the environment. So he could not claim not to have known. Hence, the vast majority of folk will say he harmed the environment intentionally (knowing fully well that he did not proclaim such an intention). Now 29% do think that he must have intended the harm (probably because they were told that he knew the harm would occur if they went ahead with the program in question). The remaining 58% used the label ‘intentional’ to blame the CEO’s actions and to discourage such activities. That is, they used the term for its pragmatic dimension, not because they theorize that one can intentionally do what one does not intend. My considered view is that most people given the survey have no very articulate theory of intentional actions, cybernetics, root origins of words or anything else. They do however understand the role of intentional talk in the praise-and-blame-game of life, and they do disapprove of harming the environment knowingly and not caring at all about it (which is just what they are told the CEO is doing in Knobe’s survey). So I believe strongly that the best

explanation of Knobe's data is that the results are due to the pragmatic dimension of intention talk.

Knobe knows the line I take on this and he and Burra make a few remarks in their paper in this journal about why they think it is wrong (before they get to their survey using Hindi language). They note that Hugh McCann (forthcoming) did a survey where he found that 80% of subjects were willing to say the chairman harmed the environment *intentionally*, while only 20% said it was the chairman's *intention* to harm the environment. This result is even more dramatic than the one found by Knobe. As McCann⁵ admits, and Knobe and Burra repeat, "people are perfectly willing to apply the word 'intentionally' in cases where they would not be willing to apply the word 'intention.'" (p.10) This is quite right, but the question is why are they willing to do so? I think it is because of the pragmatic dimension of intentional language.

Further, as I have also argued, it is harder to determine (of another) what is in one's heart when acting than in one's deed. It is easier to judge that one did something intentionally from one's actions than to judge the exact content of one's intentions. I may know by watching you that you rammed your car into the car in front of you on purpose, but it may be much harder to divine the exact content of your intention (anger, fear, suspicion). It simply may be this fact that makes subjects reticent to say what the CEO intended (thereby, driving down the number who say he intended to harm the environment...a specific purpose). Nonetheless, they may well want to blame him for

⁵ For reasons I won't go into here, McCann (forthcoming) actually agrees with me that these data do not show that one can intentionally do A without intending to do A nor that they show the folk really believe this. McCann's explanation for why these results don't show this is slightly different but friendly to mine.

harming the environment knowingly and the only choice they are offered⁶ for blaming in Knobe's or McCann's surveys is to select the option of saying the CEO harmed the environment "intentionally." That is what they choose on the survey. So the fact that it is harder to know exactly what is in the mind of an agent, harder than to know something is there guiding the agent's intentional behavior, also helps answer Knobe and Burra's question of "*why* people tend to use the word 'intentionally' (rather than 'intention') when they are engaged in discussion of praise or blame." (p. 12)

Knobe and Burra also say, in discussing the views of Adams & Steadman that we don't "provide any reason to believe that the concept of intentional action is more closely connected to the concept of intention than one might otherwise have assumed." (p. 11)

This is why I started with the cybernetic model of purposive behavior in section I. Steadman and I were all along working from that perspective. So it is quite wrong to say that we never provided any reason to connect intentions to intentional action (other than the morphologies of the English words 'intention' and 'intentional' or their common roots).

Steadman and I⁷ earlier also suggested that Knobe should have tested to see if the folk would assent to the sentence (a) "The chairman unintentionally harmed the environment" or (b) "It was not the chairman's intention to harm the environment." To use one of these sentences may have the pragmatic effect of suggesting letting the chairman off the hook (and we did not think subjects would want to do that). In the paper in this journal, Knobe and Burra complain that "the difference between these two

⁶ Annie Steadman (manuscript) and I conducted surveys of our own where we gave subjects the choice of saying that someone did act A knowingly, but not intentionally. We found that we were able to get subjects to make this distinction in some cases. But this option was not available in the surveys of Knobe or McCann.

sentences does not derive entirely from the fact that ‘unintentionally’ is an adverb and ‘intention’ is a noun. There must be some fundamental difference between the concept of intentional action and the concept of intention that accounts for people’s willingness to use sentence (b) but not sentence (a).” (p. 11)

The whole point of the pragmatic approach is to show that this very conclusion is unwarranted. What people are willing to *say* may not be a perfect guide to their *concepts of intention and intentional action*, **because** the intentional talk has a *pragmatic* as well as a *literal* dimension. When we see what people *say* or “are willing to say,” they may be willing to say this for pragmatic (not literal) reasons. Knobe’s data does not, in and of itself, tell us which use people are making of the intentional language—literal or pragmatic use.

Let me make two other points quickly before moving to the next section. Knobe and Burra mention an idea defended by Bratman (1987) and Mele (1992) that what may make an action A intentional is not that the agent intended A but that the agent intended B (not identical to A) and foresaw that B may lead to A. There are many problems with this view. One is that the argument for the need for the view is flawed, as I argue elsewhere.⁸ Bratman’s argument for the need for this view is flawed, and independently of that, the view itself is flawed. For one, how are we to tell when my intending to B (knowing A may follow) will make A intentional and when it won’t? No one has devised a good metric. Not Bratman. Not Mele. Not Knobe & Burra. To convince you that this is a problem, consider this example. I am trapped on a ledge. The only way off is to jump a chasm. I know I may not make it. I know that if I don’t make it I may fall to my

⁷ Adams & Steadman (2004a)

⁸ Adams (1986, 1997) and Adams (forthcoming).

death, but falling to my death is not my goal. I want to live. If I jump and don't make it, I jump (B) intentionally. And I foresee that B (jumping) may lead to A (my falling to my death). But I *do not* do A (if I do A), *intentionally*. That would be *suicide* (viz. doing something that intentionally leads to my death). Since I don't want commit suicide in trying to escape, there needs to be a metric to determine when doing something foreseen but not intended is done intentionally, and when it is not. I don't think such a metric can be invented. Merely foreseeing that something you intend (B: making profit) will lead to something you don't intend (A: harming the environment), does not automatically mean that you're doing A (if you do) intentionally. And that is partly why I think that the only thing that can make my doing A intentional is in part that I intend to do A.

The last point is that Knobe and Burra suggest that a plausible view (theirs) of why one finds the data one does in Knobe's surveys is that 'intentional' and 'intentionally' "express two entirely different concepts. (p. 14) Indeed they are different concepts (one of a mental state, the other of an action), but that is not what Knobe and Burra mean. They mean that these notions are *completely unrelated* (one "purely psychological" and one a matter only of "moral responsibility"). This view seems to me to be desperate and hopeless. On my own "pragmatic" view, to call an action 'intentional' can be related to moral responsibility—both literally and in the pragmatic use of that term. *But it need not*. If I say "Tom scratched his nose intentionally" what I have said may well be true. And what makes it true will be Tom's psychology (at least in part). He will have had the intention or goal to scratch his nose. Nothing much about moral responsibility hangs on this at all, yet it can be *perfectly true*. So there simply cannot be the divide that Knobe and Burra are proposing for the proper (literal) use of

these two notions (and terms). Again, please note that this has nothing to do with the morphology of the terms in English that we use to denote intentions or purposive behavior. It has everything to do with the cybernetic model of purposive behavior and what it requires.

III Linguistic Evidence from Other Cultures

Knobe and Burra build their paper around what I have argued is a mistaken premise, viz. “It is simply taken as a given that since, since the English word ‘intentional’ sounds so much like the English word ‘intention,’ there must be a tight connection between the concepts these words express.” (p.3). It is taken as a given by whom? In Section I, I made no such appeal, nor do McCann or any others who think as I do appeal to sounds of words or roots of words or morphology to defend our views. Knobe and Burra deny that such linguistic accidents matter. So do I.

Next they look at Hindi and the judgments of Hindi speakers with respect to actions where they would use the words *jaan* (‘know’) and *jaan-bujkar* (which they translate as ‘intentional’). Using the same vignettes where they ask about the CEO “helping” or “harming” the environment, surveyed Hindi speakers said the following: 90% said the CEO *jaan* (knew) he would help the environment but only 14% said he helped it *jaan-bujhkar* (intentionally). In the harm condition, 80% said he *jaan* (knew) he would harm the environment, and a full 75% said he harmed it *jaan-bujhkar* (intentionally).

To my mind this is quite unsurprising. It is exactly what would be expected on the pragmatic view that I discussed above. Hindi speakers want to discourage one for harming the environment “intentionally” (or knowingly, given the root derivation of

'jaan-bujhkar' from *'jaan'*). They know the best way to do this is to use the pragmatic weight of the term *'jaan-bujhkar.'* Similarly, they don't want to praise him for helping the environment knowingly, since he says (in Hindi) the equivalent of "I don't care at all about helping the environment." So only 14% say he helped *'jaan-bujhkar.'*

I would think it is something of an embarrassment to Knobe and Burra that this is so. For they say that they find little "evidence that the concept of acting intentionally is more closely related to the concept of intention than it is to the concepts of wanting, trying, foreseeing, and so forth" (p.21). If true, the subjects knew that the CEO could *foresee* that his actions would help the environment. He was told this by his advisors. But subjects could also see that he *didn't* intend the help. So if Knobe and Burra are right, and *'jaan-bujhkar'* works more like "foresee" (in this case *'jaan'*), than 'intend', then *why didn't more* subjects think the CEO helped *'jaan-bujhkar'* (knowingly)? My answer is that even though they knew he could foresee the help, they didn't want to praise him so they didn't say his action was done *'jaan-bujhkar'* (intentionally). This is just as the pragmatic view predicts, but it is **not** what is predicted if *'jaan-bujhkar'* acts more like the English word "knowingly" than like the English word "intended."

To be fair, Knobe and Burra do see this fact about their data, but all they say is that "this second finding indicates that the word *'jaan-bujhkar'* is not simply an adverbial form of *'jaan.'* The problem for them then is that if it is acting more like 'intentionally' than 'knowingly,' then they must retract their claim that *'jaan-bujhkar'* acts more like "foresee" than "intend." They can't have it both ways.

Finally, it seems that they are trying to have it both ways when they say about "intentionally" and *'jaan-bujhkar'* that "despite their radically different morphologies,

these two words are expressing the same basic concept.” This cannot be correct, for subjects don’t judge this way in the “help” condition, *as the data plainly shows*.

Now Knobe and Burra will say, indeed they do say, that when ‘jaan-bujhkar’ acts like “knowingly” and when it acts like “intentionally” all “depends ... on the *moral* status of the behavior—so that it is closer to that of ‘intention’ for morally good behaviors and closer to ‘jaan’ for morally bad behaviors” (p. 21). So indeed they are trying to have it both ways! To my mind this means that they are saying that the literal meaning of these words must change when we apply them to actions judged morally good *versus* when applied to actions judged morally bad. This would literally mean that the words acted like indexicals (“I” and “here” and “now”) or words of ambiguity (river “bank” *versus* financial “bank”).

I ask you: what would be more surprising? That the meaning of ‘intentional’ or ‘jaan-bujhkar’ literally changed (truth conditions actually changed) when applied to morally good actions *versus* when applied to morally bad actions. Or would it be more surprising that the literal meanings of the words remained the same, but the pragmatic use of the words (to blame or to discourage others from harming the environment) was at work in the minds of the subjects when using them to praise or blame? From my perspective, the former would be much more surprising than the latter, especially since there is no other reason to think that the meanings of these words change or are ambiguous or act like indexicals. If the meanings did change, then the word ‘intentionally’ in “Tom intentionally (‘jaan-bujhkar’) scratched his nose” would literally

have a different meaning than in “Tom intentionally kicked his innocent child.” If you accept the latter, believe Knobe and Burra. If you don’t accept the latter, believe me.⁹

IV CONCLUSION

In this paper I have given reasons to think that intentions confer intentionality or purpose upon our actions. The reasons I have given have nothing to do with the root derivation or morphology of terms ‘intention’ and ‘intentional’ in English (contrary to suggestions of Knobe & Burra).

I have also discussed the claims of Knobe and Burra that the concepts *intention* and *intentional* are completely different and apply to different things for different reasons. I’ve suggested that this is not true, and that if it were true, it would be a problem for their explanation of their own data. To the contrary, I have explained why the two concepts are related for the reason that the former is the concept of a mental state that confers purpose upon behavior turning it into the latter. This is a matter of the nature of things, not of words.

I have explained why I think the *pragmatic* account of the use of intentional language that I have proposed is the best way to explain the very intriguing data that Knobe and Burra have uncovered. This account explains the asymmetry of judgment in the case where the CEO causes harm (rather than help) to the environment—and why subjects say the first (but not the second) was intentional. The say it was intentional

⁹ Sachie Kotani and Yurie Hara, both native Japanese speakers, tell me that in Japanese ‘ITO’ means ‘intention’ and ‘ITO-TEKI(NA)’ means ‘intentional.’ So they have the same root. Still, ‘ITO-TEKI(NA)’ is usually used for a negative/bad action and would not work well for application to a good action. So even in languages where the terms for intentions and for intentional actions *have the same root*, there can be pragmatic or conventional restrictions on use. For the term ‘WAZATO’ meaning ‘on purpose,’ this term too is typically used for doing negative or bad actions on purpose and would be strained if used for good

because they are employing the pragmatic effect of the *saying*, not the literal meaning of *words said*.

Finally, I have explained why I do not find any news in the data of Knobe & Burra using words from Hindi. This data is consistent with the data in English and to my mind the pragmatic account will explain it as well. I have also alluded to the fact that I find it much more plausible to account for the data Knobe and Burra give (both in English and in Hindi) on the pragmatic account rather than think the actual *literal meanings* of the words ‘intentional’ or ‘jaan-bujhkar’ change when applied to good (or morally neutral) actions *versus* when applied to morally bad actions. There is no evidence or other reasons (independent of Knobe & Burra’s surveys) to think this happens with these particular terms. There is evidence that people employ these terms with both their *literal* and sometimes with their *pragmatic* uses. So I find the pragmatic account to be much more plausible in explaining the available data from surveys, and I strongly recommend this account to you.

References

Adams, F. (forthcoming). Trying With the Hope.

Adams, F. (1986). Intention and Intentional Action: The Simple View, *Mind & Language*, 1, 281-301.

Adams, F. (1997). A Cognitive Theory of Trying, In G. Holmstrom-Hintikka & R. Tuomela (Eds.) *Contemporary Action Theory*, Vol. I, 287-314, Dordrecht: Kluwer.

actions. There can be many cultural and pragmatic reasons for these restrictions that go beyond the literal truth conditions or meaning of terms.

- Adams, F. (2003). The Informational Turn in Philosophy, *Minds and Machines*, 13, 471-501.
- Adams, F. & Steadman, A. Manuscript. Folk Concepts, Surveys, and Intentional Action.
- Adams, F. & Steadman, A. 2004a: Intentional Action in Ordinary Language: Core Concept or Pragmatic Understanding?, *Analysis*, 64, 173-181.
- Adams, F. & Steadman, A. 2004b: Intentional Action and Moral Considerations: Still Pragmatic, *Analysis*, 64, 268-276.
- Knobe, J. (2003a). Intentional Action in Folk Psychology: An Experimental Investigation, *Philosophical Psychology*, 16, 309-324.
- Knobe, J. (2003b). Intentional Action and Side Effects in Ordinary Language, *Analysis* 63, 190-193.
- Knobe, J. (2004). Intention, Intentional Action and Moral Considerations, *Analysis*, 64, 191-187.
- Malle, B. & Knobe, J. 1997. The Folk Concept of Intentionality, *Journal of Experimental Social Psychology* 33: 101-121.
- McCann, H. forthcoming: Intentional Action and Intending: Recent Empirical Studies, *Philosophical Psychology*.