

APPENDICES

APPENDIX A.

THE ASYMMETRY THESIS

J.M. Fischer has argued¹ that there is an important asymmetry in responsibility ascriptions between actions and omissions. Fischer offers two kinds of cases which are to show that we indeed do harbor intuitions supporting the thesis. We will see that if one casts Fischer's cases in our apparatus, one will be able to explain the intuitions without needing to postulate any asymmetry between action and omissions. Indeed, the apparatus handles objections put forward against the asymmetry thesis equally well.

1. The Asymmetry Thesis

Fischer agrees with Frankfurt that responsibility for actions does not depend on our ability to have done otherwise. He suggests, however, that responsibility for omissions does depend on our ability to have done otherwise. This is the asymmetry thesis: there is an asymmetry in responsibility ascriptions between actions and omissions.

To substantiate the thesis Fischer considers two kinds of cases: Frankfurt-type cases of actions for which we are responsible despite not having been able to do

¹ The original thesis appeared in John Martin Fischer's "Responsibility and Failure," *Proceedings of the Aristotelian Society* 86 (1985/86), 251-270. It has been elaborated in a paper with M. Ravizza ("Responsibility and Inevitability," *Ethics* 101, 1991, 258-278), and is upheld by Fischer in *The Metaphysics of Free Will. An Essay on Control* (Oxford: Basil Blackwell, 1994). In the meantime, the thesis has received some attention, see Randolph Clarke, "Ability and Responsibility for Omissions," *Philosophical Studies* 73 (1994), 195-208; Harry Frankfurt, "An Alleged Asymmetry between Actions and Omissions," *Ethics* 104 (1994), 620-623; Ishtiyaque Haji, "A Riddle Regarding Omissions," *Canadian Journal of Philosophy* 22 (1992), 485-502; Alison McIntyre, "Compatibilists Could Have Done Otherwise: Responsibility and Negative Agency," *Philosophical Review* 103 (1994), 453-488; David Zimmerman, "Acts, Omissions and 'Semi-Compatibilism'," *Philosophical Studies* 73 (1994), 209-223.

otherwise, and then cases of omissions for which we are not responsible *because* we have not been able to do otherwise. Let us consider them in turn.

Frankfurt-Type Cases purport to illustrate that there are situations where we would hold the agent responsible despite the fact that he could not have done otherwise. Consider the following case: Jones decides to kill the mayor of the town. He carries out his plan to the letter, shoots the mayor who subsequently dies. Unbeknownst to Jones, evil scientists have implanted a device into Jones' brain which, were Jones to decide not to kill the mayor (or waver after his decision) would have swayed Jones to kill the mayor anyway. The intuitions about cases of this sort have been almost uniform. Jones is responsible for killing the mayor. At the same time, it has been claimed, Jones could not have done otherwise: he could not have not killed the mayor.

Fischer's Cases of Omissions. Here is a case of an omission for which, Fischer suggests, the agent is not responsible. Jones does not have any fancy mechanism in his brain. He is strolling along the beach when he sees a child struggling in the water. Though he believes he can rescue the child with little effort, he decides not to go to the trouble. The child drowns. Unbeknownst to Jones, had he jumped into the water, the sharks patrolling the beach would have attacked him, so Jones could not have saved the child after all.

Fischer believes that this is a case where Jones is not responsible for failing to rescue the child precisely because he could not have rescued her. Fischer does not deny that Jones is responsible for something. He is responsible for his "failure to *try* to save the child (and his failure to jump into the water, etc.)."² But he is not responsible for his failure to save the child.

Assuming that there are no qualms with respect to the intuitions themselves, there is indeed a striking difference between these cases. In the case of the action, we are inclined to think that the agent is responsible for it. In the case of the omission, we are inclined to think that the agent is not responsible for it. Yet both cases are similar with respect to the fact that the agent could not have done otherwise, a fact that has traditionally been held to be of great significance in ascribing responsibility. It would

² J.M. Fischer, "Responsibility and Failure," *op. cit.*, p. 253.

indeed be plausible to agree that this indicates that there is an asymmetry in the responsibility conditions for actions and omissions if one agreed that the cases are similar with respect to the responsibility-engendering condition. Here in outline is the argument that they are not.

I will assume (but not defend the assumption) that the fact that a performance is something the agent has done (in the sense discussed in Chapter VI, section 3) is a necessary but not a sufficient condition for the agent's being held responsible for the action (under that description). In other words, if an agent is held responsible for an action under a description, it follows that it would have been reasonable_A to expect of the agent that she perform the action under that description.³ I will first argue (section 2) that the cases Fischer takes to support the asymmetry thesis are dissimilar in this respect. In the Frankfurt-type actions, we hold the agent responsible for ϕ ing and ϕ ing counts as something the agent has done: it is reasonable_A to expect of her that she ϕ . In the Fischer-type omissions, we do not hold the agent responsible for ϕ ing and ϕ ing counts as something the agent happens to do: it is unreasonable_A to expect of her that she ϕ . The fact that the cases are dissimilar in what I assume to be the responsibility-engendering condition (that what the agent is held responsible for is something she has done rather than happened to do) does not yet disprove Fischer's asymmetry thesis. It would if Fischer had to agree with my assumption but, as indicated, I do not offer a defense of it (see footnote 3). What this will demonstrate is the fact that there is an alternative explanation of our intuitions concerning the cases that appeals to the assumption and which does not demand that we postulate an asymmetry between actions and omissions. In support of this alternative explanation I then consider (section 3) an omission which

³ This thought might indeed be the healthy core of what Mackie has called the "straight rule of responsibility" according to which we are responsible for all and only intentional actions (*Ethics. Inventing Right and Wrong* [New York: Penguin Books, 1977]). The rule is too restrictive. We are often held responsible for unintended consequences of our actions, for unintentional omissions, etc. It does not seem implausible to suggest that the notion of it being reasonable_A to expect something of an agent, which as I argued in Chapter VI can replace the notion of a performance being intentional in the understanding of the nature of action, could also replace the latter in the understanding of the performance for which we are responsible. This is a conjecture that needs developing. In particular, I have not offered a systematic treatment of consequences of actions, which would be required before any such hypothesis can claim to be more than a conjecture.

has the structure of a Frankfurt-type action and a case of an action which has the structure of a Fischer-type omission. Indeed, we shall see that our intuitions coincide with the suggested explanation thus vindicating the assumption.

2. The Reconstruction of the Two Types of Cases

Let us compare the two kinds of cases in the apparatus developed.

Frankfurt's Cases. Does the presence of the counterfactual intervener render it unreasonable_A to expect of Jones that he kill the mayor? One might think that it does. After all, given the presence of the counterfactual intervener it is determined that the mayor will die at Jones' hands. It would thus seem that the presence of the counterfactual intervener is systematically correlated with the pf-fulfillment of the expectation that Jones kill the mayor. However, I have argued that the presence of the counterfactual intervener violates Principle III (p. 112). There are exactly two avenues to the mayor's death at Jones' hands envisaged in the example. First, Jones might decide to kill the mayor and so kill him, in which case we are invited to suppose that his decision to kill the mayor (*K*) is systematically correlated with the (agentive) fulfillment of the expectation that he kill the mayor. Second, Jones might not decide to kill the mayor in which case the counterfactual intervener will take over leading Jones to kill the mayor. In this case, we are invited to suppose that the scientist's intervention (*C*) is systematically correlated with the (non-agentive) fulfillment of the expectation that he kill the mayor.

The reason why we are originally inclined to think that the presence of the counterfactual intervener would be systematically correlated with the fulfillment of the expectation that Jones kill the mayor relies solely on the fact that the case is constructed in such a way that either one or the other condition operates. In other words, given the details of the case, we might be tempted to construe condition *K-or-C* as a defeating condition. Principle III blocks this move. In view of the fact that we already understand the operation of the existing conditions (here *K* and *C*) the "new" condition *K-or-C* does not defeat the reasonableness_A of the expectation that Jones kill the mayor. It is thus reasonable_A to expect of Jones that he kill the mayor despite the presence of the counterfactual intervener.

Given the account of Chapter VI, Jones's following through his decision to kill the mayor is something he has done rather than something he happened to do. Given that it is reasonable_A to expect of Jones that he kill the mayor, that the mayor does get killed at Jones' hands (and that the counterfactual intervener does not *actually* intervene), Jones performed an action of killing the mayor. (If the counterfactual intervener did actually intervene leading Jones to the killing, the intervention would render the expectation unreasonable_A, and so we could at most say that his killing the mayor is something he happened to do.⁴)

Fischer's Cases of Omissions have a different structure. Here the potential defeating condition with respect to the expectation to save the child, the presence of the sharks, does not counterfactually depend on the agent's decision. It is a condition that operates up front as it were.

In view of the sharks patrolling the beach, it would be unreasonable_A to expect of Jones that he save the child. This is because the presence of the sharks is assumed to be systematically correlated with the pf-frustration of the expectation to save the child. (We are asked to suppose that it is in fact impossible for Jones to do so — he would be attacked before he ever got to the child.)

Is it reasonable_A to expect of Jones that he not save the child? Once again, the answer is negative. The presence of the sharks guarantees that the any expectation not to save the child will be pf-fulfilled. The presence of the sharks is a defeating condition of the second kind with respect to the expectation not to save the child (and of the first kind with respect to the expectation to save the child).

As long as it would be unreasonable_A to expect of Jones that he prevent the sharks from attacking, it is unreasonable_A to expect of Jones that he rescue the child as well as that he not rescue the child. Hence, in view of Chapter VI, Jones' failure to save the child is not something he does.⁵

⁴ In fact, in this case the defeating condition is severe enough to make it unreasonable_A to expect the performance of the agent under all descriptions.

⁵ This echoes Frankfurt's response to the case: "The real reason [why Jones bears no moral responsibility] is that what he does has no bearing at all upon whether the child is saved. The sharks operate both in the

In both these cases, we have arrived at the right kind of judgment about them without presupposing that there is a deep asymmetry between actions and omissions. The difference concerns rather the structure of the cases. In Fischer's cases of omissions, there is a defeating condition which makes it unreasonable_A to expect of the agent that he performs the action as well as that he does not perform it. In the Frankfurt-type cases, on the other hand, the presence of the counterfactual intervener does not count as a defeating condition: only the actual intervention by the scientist would count as such. Since in the actual situation, as it is construed in the Frankfurt-type case, the only defeating condition (the scientist's intervention) does not occur, it is reasonable_A to expect of Jones that he kill the mayor.

3. Frankfurt-Type Omissions and Fischer-Type Actions

The case against the asymmetry between actions and omissions can be strengthened further if we could find examples of Frankfurt-type omissions, for which we would be responsible, and examples of Fischer-type actions, for which we would not. Such examples can indeed be found.

Frankfurt-Type Omissions. Let me begin by illustrating an example of omission that exactly parallels the structure of Frankfurt-type actions.⁶

Brown has an implant similar to Jones'. She is walking along the beach and sees a child struggling in the water. Though Brown cannot swim, she can throw a life jacket but decides not to. The child drowns. Unbeknownst to Brown, had she shown any inclination to try to save the child, the implant would have been activated as a result of which Brown could not attempt to rescue the child after all. As it happens, the implant did not need to be activated. In this case, the intuition seems to be that Brown is morally responsible for failing to throw the jacket to the child even though she could not have done otherwise.

actual and in the alternative sequences, and they see to it that the child drowns no matter what John does" ("An Alleged Asymmetry between Actions and Omissions," *op. cit.*, p. 623).

The case is exactly parallel to Frankfurt-type actions. It might appear that it is unreasonable_A to expect of Brown that she not throw the life jacket, for given the arrangement of the case, the expectation not to throw the life jacket will be systematically pf-fulfilled. However, as before in the Frankfurt-type case, there is in fact no defeating condition at work. The alleged defeating condition is a composite of (a) the decision not to throw the life jacket *D* which is supposed to be systematically correlated with the (agentive) fulfillment of the expectation not to throw the jacket, and (b) the scientist's possible intervention *C* which is systematically correlated with the (non-agentive) fulfillment of that expectation. It is because the case is so constructed that either *D* or *C* will occur that we might think a defeating condition is in place. In virtue of Principle III, however, *D-or-C* does not qualify as a defeating condition. Hence, despite *D-or-C* it is reasonable_A to expect of Brown that she not throw the life jacket. By similar reasoning (which exactly parallels the Frankfurt-type case), despite *D-or-C* it is reasonable_A to expect of Brown that she throw the life jacket. The situation would change were the scientist to intervene. The scientist's actual intervention would qualify as a defeating condition. It would no longer be reasonable_A to expect of Brown that she throw the life jacket if the scientist intervened.

Because in the actual situation, the scientist does not intervene, it is reasonable_A to expect of Brown that she not throw the life jacket. So, when in the actual case, she does not throw the life jacket, while the counterfactual intervener does not intervene, her not throwing the jacket is something she does.

Fischer-Type Actions. It is in general more easy to describe a Fischer-type omission than action but perhaps the following example will bring the point home. It does not involve a counterfactual intervener. Smith wants to switch on the light. He presses the switch. The light comes on. It might appear that Smith switched on the light. However, unbeknownst to Smith, the light would have come on at exactly the moment that Smith actually pressed the switch, but independently of Smith's intervention. It seems intuitive to describe the case as that of Smith having nothing to do with the light going on. (Indeed Smith could

⁶ The case is borrowed from I. Haji, "A Riddle Regarding Omissions," *op. cit.* A similar case is

not have done otherwise: he could not have not switched the light.) It would be inappropriate to hold Smith responsible for switching on the light. Smith might still be held responsible for flipping the switch, but not for actually switching the light on. This is indeed borne out if we ask whether it was reasonable_A to expect of Smith that he switch on the light.

Given that the light will come on at t , is it reasonable_A to expect of Smith that he not switch on the light at t ? It seems clear that the expectation would be unreasonable_A. The fact that light will come on at t is systematically correlated with the pf-frustration of the expectation that Smith not switch on the light at t and with the pf-fulfillment of the expectation that Smith switch on the light. Hence, it was unreasonable_A to expect of Smith that he switch on the light at t .

4. Final Remarks

The asymmetry thesis concerns the asymmetry in ascriptions of responsibility. I have not defended any view regarding the conditions of responsibility. I have rather followed a simpler strategy. I have assumed that it is necessary for an agent's being responsible for ϕ ing that ϕ ing count as something the agent has done (i.e. that it was reasonable_A to expect of the agent that she ϕ). I have then shown that the assumption allows us to understand the difference in our dispositions to hold the agent responsible in the cases of Frankfurt-type actions and Fischer-type omissions (section 2). If the assumption is correct this has nothing to do with the fact that the former are actions and the latter are omissions, but rather with the fact that in the former cases we can say that the performance is something the agent does while in the latter it is only something the agent happens to do. I have then vindicated the suggestion by showing that we would be inclined to hold the agent responsible in the case of a Frankfurt-type omission, while we would not hold the agent responsible in the case of a Frankfurt-type action (section 3). This supports the view that there is no fundamental difference between actions and omissions with respect to responsibility ascriptions.

APPENDIX B.

ACTION AS A PERFORMANCE INTENTIONAL UNDER A DESCRIPTION

The currently most popular answer to the question whether anything has been done appeals to what might be called the criterion of intentionalness. The criterion was first proposed by Anscombe¹ and later adopted by Davidson²:

- (I) An event e is an action if and only if e is intentional under some description.

I will argue that while (I) might be useful for those theorists of action who aim to understand the category of intentional movements, it must be rejected by anyone aiming to understand the category of action as a unit of conduct.

I begin by sketching six distinct ways in which (I) may be understood (section 1). In section 2, I discuss the ramifications of the most plausible (non-reductive explicatory) reading of (I). In section 3, I argue that the considerations raised about the non-reductive explicatory reading of (I) actually constitute a reason to take the reductive explicatory readings of (I) as being either circular or faulty.

1. A Methodological Prelude

The first point that ought to be raised about (I) concerns its status. (I) might be taken to constitute an *analysis* either of the concept of action or of the concept of being intentional under a description. Alternatively, it could be thought of as not analyzing but rather as *reporting a conceptual connection*. It seems undeniable that Anscombe does

¹ *Intention* (Ithaca: Cornell University Press, 1957).

² "Agency," in *Essays on Actions and Events* (Oxford: Clarendon Press, 1980), pp. 43-61.

not intend (I) as a reductive analysis. It is deniable, though there are good reasons not to deny, that Davidson likewise does not intend (I) to be reductive.³

However, both the reductive and the non-reductive interpretation of (I) can come in a number of flavors. On the reductive side, (I) functions as something of a definition and can correspondingly be understood in at least three ways: as a conventional definition (RC), as an analytic definition (RA) and as an explication (RE).

(RC) It may be understood as *stipulating* that one concept is to be understood in terms of the other. In this case, the definition is purely conventional, and it does not depend on any prior usage of the definiendum. Such a definition cannot be criticized on cognitive grounds, for it coins a new concept. It seems rather clear that (I) is neither meant nor taken as a conventional definition.

(RA) It may be understood as an *analytic definition*. In this case, the concepts are supposed to be well-established with clear areas of application. The purpose of the analytic definition is to analyze one concept in terms of the other; it is to suppose that the definiendum can be understood in terms of the definiens, where the definiens is treated as logically prior to the definiendum.⁴ In general, an analytic definition is subject to criticism if the extensions (or the intensions) of the concepts are not identical, as well as if the areas of vagueness or imprecision do not overlap. An example of an analytic definition could be⁵ “A bachelor is a married man,” or “Mental states are (nothing but) physical states.”

(RE) Finally, and most plausibly, (I) may be understood as an *explication*. In such a case, the explicandum is treated as a concept whose intension or extension is sharpened, clarified or illuminated by explicating it in terms of the explicans. Unlike an analytic definition, an explication requires neither that the areas of

³ See Donald Davidson, “Freedom to Act,” in *Essays on Actions and Events*, *op. cit.*, pp. 63-81.

⁴ It is hard to escape the impression that analytic definitions are also in some sense conventional. After all, if the two concepts are well-established, then the claim that one ought to be understood in terms of the other rather than the other way around seems purely conventional or arbitrary. It is only if they are embedded in a larger explicatory project that they may be useful. In such a case, the ordering of concepts is set by the explicatory definitions.

⁵ “Could be” because the statement is probably better construed as reporting a pre-existing conceptual connection, see (NA) below. It would be an analytic definition if one were to treat the concepts of being married and being a man as somehow more basic than the concept of a bachelor (see footnote 4).

application of both concepts overlap completely nor that their intensions match. In this respect, an explication is partly stipulatory. Unlike a conventional definition, however, an explication does rely on some overlap between extensions and intensions of the concepts. In this respect, an explication is partly responsive to existing conceptual connections. As such an explication is criticizable, although its value is supposed to lie in how well it functions within a system of explications, in the extent to which it introduces conceptual order.

A famous example of an explication is the statement “Water is H₂O.” Unlike a conventional definition, it does rely on preexisting conceptual connections. Some of the conceptual connections involving “Water” are preserved when “H₂O” is substituted (providing other substitutions are made). Unlike an analytic definition, there is no pretense that the intension and extension of the concepts is the same, so that many of the conceptual connections are either rejected or stand in need of explanation (e.g., “Water is liquid,” “Water does not break windows though ice sometimes does,” “The water in this lake contains all kinds of dirt and chemical substances”).

What all these types of definition have in common is that the definiendum is taken to be of a different logical order than the definiens. But they differ in the extent to which the definition is responsive to preexisting conceptual connections. A reductive-conventional statement does not require any pre-existing conceptual connections. A reductive-analytic statement requires that the conceptual connections overlap entirely. And a reductive-explicatory statement requires a partial overlap.

Construing (I) non-reductively means that we must abandon the idea that one concept is somehow logically prior to another. But a non-reductive reading likewise may be responsive to preexisting conceptual connections to different degrees: it may *establish* a conceptual connection (non-reductive-conventional; NC), it may *report* a conceptual connection (non-reductive-analytic; NA) or it may *partly report* and *partly establish* a conceptual connection (non-reductive-explicatory; NE).

(NC) It may be understood as stipulating that two concepts are to be understood in terms of one another. In this case, the claim is purely conventional. It does not

depend on any prior usage of the concepts. Such a claim cannot be criticized on cognitive grounds, for it coins the conceptual connection.

(NA) It may be understood as reporting a pre-existing conceptual connection. In this case, the concepts are supposed to be well-established with clear areas of application. The purpose of the claim is to report the connection existing between the concepts. If there are any areas of vagueness they overlap.

An example of a statement reporting a conceptual connection is “A bachelor is a married man.”

(NE) It may be understood as partly reporting but partly establishing a conceptual connection. This is most likely when the areas of indeterminacy or vagueness of the two concepts do not overlap, so that in some cases, one of the concepts can clarify the other, while in other cases, the clarificatory roles are reversed.

Again, what is common to all these options is that the concepts involved are treated as being of the same order. However, they differ in the extent to which they are responsive to existing connections. A non-reductive-conventional statement does not require any pre-existing conceptual connections. A non-reductive-analytic statement requires that the conceptual connections overlap entirely. And a non-reductive-explicatory statement requires a partial overlap.

2. Ramifications of a Non-Reductive (NE) Reading of (I)

If we give (I) a non-reductive reading, it seems most plausible to construe it as explicatory (NE). It seems obvious that (I) is not intended to *establish* a conceptual connection in the manner of (NC). It is less obvious that (I) is not well construed as simply reporting a conceptual connection as in (NA). In order to show this, it will pay to look at cases where it is plausible to suppose that one concept helps to clarify the other. Among the cases where the concept of being intentional under a description helps our intuitions about the concept of action are cases of negative actions. Among the cases where the concept of action helps our intuitions about the concept of being intentional under a description are cases of spontaneous actions (which philosophers sometimes cast as actions performed only with an intention-in-action, not on prior intention). If so, then to the extent that we want to uphold (I), we ought to recognize it as not simply reporting a

conceptual connection between two concepts, but rather as partly reporting and partly establishing it. Let us see that this is so.

Let us take the class of cases where the concept of action is sharpened by our grasp of the concept of being intentional under a description. The most important cases here are those of negative actions. Prima facie we might think that there are a number of classes of negative actions. Although the following list is not exhaustive it should be sufficiently suggestive. (a) An agent is tempted to eat yet another cookie but has firmly resolved to go on a strict diet. His not eating the cookie (refraining from eating it) is something he does. (b) An agent is obligated to file a report but decides not to do so. In this case his not doing it (intentionally omitting to do it) is something he does. (c) An agent accompanies his friend to a party, where another person attacks his friend in conversation. Our agent does absolutely nothing but without intending to do anything either. He just idly stands there. His failing to come to his friend's help is also something he does (something that his friend will rightfully blame him for). (d) A person oversleeps as a result of which he fails to come to a crucial meeting. Once again, one might think it something the agent does.

Our intuitions regarding cases (a) through (d) are not uniform. It suffices for my purposes here to demonstrate that there are some among us whose intuitions favor the inclusion of all these cases under our actions.⁶ There are others, however, who favor the inclusion only of cases (a) and (b). What distinguishes cases (a) and (b) from (c) and (d) is precisely the fact that (a) and (b) are intentional under the negative description of the action, while (c) and (d) are not.⁷ Indeed, some authors appeal to the fact that (c) and (d)

⁶ John C. Hall, "Acts and Omissions," *The Philosophical Quarterly* 39 (1989), 399-408; H.L.A. Hart, *Punishment and Responsibility* (Oxford: Oxford University Press, 1968); Steven Lee, "Omissions," *Southern Journal of Philosophy* 16 (1978), 339-354; Patricia G. Smith (Milanich), "Allowing, Refraining, and Failing. The Structure of Omissions," *Philosophical Studies* 45 (1984), 57-67; "Contemplating Failure: The Importance of Unconscious Omission," *Philosophical Studies* 59 (1990), 159-176.

⁷ Bruce Vermazen has argued that it is a special feature of negative actions that they need to be intentional (and so intentional under the negative descriptions) on pain of including too many negative actions. If one allowed unintentional negative actions to count as negative actions, i.e. if one allowed that as long as a performance is intentional under some description, it is a negative action under all negative descriptions, the list of negative actions would be endless. In Vermazen's words: "Certainly we don't want to say that a person is not- ψ -ing just in case he is not ψ -ing. ... It won't help much to add the rider 'if the agent is doing something' to this last, since the agent will then be doing far too many negative acts: Andy, as he sits twisting his buttons, would also be not-sweeping the table clear of canapés, not-preparing for a Channel

are not intentional under any description to dismiss any intuitions to the effect that they ought to be classified as actions.

The second class consists of cases where the concept of action has a firmer grip and helps us with the concept of being intentional under a description. It comprises spontaneous voluntary actions done with no reason, ones that it is reasonable_A though not reasonable_N to expect of the agent (see p. 152). Consider a case mentioned by Michael Bratman of spontaneous voluntary action:

Suppose you unexpectedly throw a ball to me and I spontaneously reach up and catch it. My catching it is under my control and voluntary; it is not just like a mere reflex blinking of my eye. But my action is relatively automatic and unreflective, so it may seem strained to suppose that its etiology must involve a distinctive attitude of intending, given that we are understanding intending largely in terms of its role in planning.⁸

(What is noteworthy here is Bratman's simultaneous appeal to the concept of "being under the agent's control," "being voluntary" and the contrast with its *not* being like a reflex action. This is exactly the path we ought to follow if we were to determine the applicability of the concept of action in terms of the absence of defeating conditions. The performance is an action as long as it is unlike performances that happen in the wrong kind of circumstances (in this case: that result from the operation of a reflex).)

One might worry here that Bratman has a vested interest in applying the concept of being intentional under a description to coincide with his concept of intention, which, as he admits, is shaped by his planning theory.⁹ He suggests in effect that cases of this sort ought to be described as "voluntary but neither intentional nor unintentional."¹⁰ So, one could argue that on an alternative understanding of intention and of the concept of being intentional under a description, we would have no problem in qualifying this case as falling right under it. But even if we ignore the virtues of Bratman's account of intention, still the point is that there *are* intuitions on which it is not obvious that

swim, not-attempting to cross the Sino-Soviet border, and so on." ("Negative Acts," in (eds.) Bruce Vermazen, Merrill B. Hintikka, *Essays on Davidson* [Oxford: Clarendon, 1985], p. 96).

⁸ Michael E. Bratman, "Moore on Intention and Volition," *The University of Pennsylvania Law Review* 142 (1994), p. 1712.

⁹ Michael E. Bratman, *Intention, Plans, and Practical Reason* (Cambridge, MA: Harvard University Press, 1987).

spontaneous actions are intentional under some description. In order to uphold (I), one must dispense with such intuitions, in other words, one must “argue” that we need to extend the concept of being intentional under a description because such cases would not be included as actions, and that they are actions seems intuitively indisputable.

In summary, I have argued that if (I) were to be construed non-reductively, it is most plausible to understand it as partly reporting but partly tightening the conceptual connection between the concepts. Cases of negative actions are ones where our understanding of what it is to be an action is sharpened by appeal to the concept of being intentional under a description. Cases of spontaneous actions are ones where our understanding of what it is to be intentional under a description are sharpened by our understanding of what it is to be an action. In other words, (I) is best understood as a non-reductive explicatory (NE) claim, but neither as a non-reductive conventional (NC) nor as a non-reductive analytic (NA) claim.

3. Circularity or Inadequacy of a Reductive Reading of (I)

It seems clear that if one wanted to construe (I) reductively, the very same sorts of cases discussed above likewise tell for construing it as an explication of either one or of the other concepts. If one were to take (I) as analyzing the concept of action in terms of the concept of being intentional under a description, the cases of mistakes or spontaneous actions would be problematic, for the concept of being intentional under a description does not straightforwardly apply to them. If one were to take (I) as analyzing the concept of being intentional under some description in terms the concept of action, one may object in a similar way that the concept of action is not firm enough with respect to negative actions.

In fact, if the considerations that I invoked against taking (I) as merely reporting a conceptual connection, and in favor of taking (I) as being non-reductively explicatory, are sound, one may argue against any attempt to take (I) to be reductive. For to the extent that it is true that our grip on the concept of being intentional under a description is sharpened by appeal to the concept of action in cases of spontaneous actions, taking (I) as

¹⁰ M.E. Bratman, “Moore on Intention and Volition,” *op. cit.*, p. 1712.

analyzing (or even as explicating) the concept of action is either question-begging or fails to capture those cases as actions. If one takes the concept of being intentional under a description to apply to spontaneous actions, its use as an explicans of action is circular. For the concept of action is crucially involved in our having a firm enough grip on the concept of being intentional under a description to apply to those cases (see above). Alternatively, if one takes the concept of being intentional under a description to apply only to cases where we have a firm enough grip on that concept (without any enrichment from the connection with the concept of action), the explication will fall short of capturing the concept of action, for it will not straightforwardly apply to spontaneous actions. Thus the explication of the concept of action in terms of the concept of being intentional under a description is either circular or faulty.¹¹

Another way of putting the point concerns the “history” of the notion of intention-in-action. Cases of actions that are not done on a prior intention are usually supposed to involve an intention-in-action.¹² If one takes the concept to be involved in a reductive analysis of the concept of action then its status will seem rather peculiar. It is most plausible to suppose that the concept of a performance being intentional under a description applies paradigmatically to cases of actions done on a prior intention, perhaps preceded by a distinct stage of deliberation. If the concept of being intentional under a description were to be limited in application to this class of cases then the class of spontaneous actions, not done a prior intention, would not be covered by it. In order for the concept of action to be properly delimited then one needed to stipulate another notion of intention, which is present even if there is no prior intention on which the agent acts, intention-in-action. Clearly, however, such an invocation of the concept renders its use in the explication circular.

¹¹ These claims are made in abstraction from any further attempts to explicate the concept of being intentional under a description. However, in view of the fact that the literature is full of controversy regarding the concept of intentional action, we might think it pretty safe to say that the concept of being intentional under a description or any of its potential explicanda are far from being precisely settled.

¹² The term is first introduced by G.E.M. Anscombe in *Intention*, *op. cit.* However, it has since acquired very different interpretations including teleological (George M. Wilson, *The Intentionality of Human Action* [Stanford: Stanford University Press, 1989]), causal (John R. Searle, *Intentionality. An Essay in the Philosophy of Mind* [Cambridge: Cambridge University Press, 1983]) and social-normative (Robert Brandom, *Making It Explicit* [Cambridge: Harvard University Press, 1994]).

Similarly, it could be argued that if (I) were taken to be an explication of the concept of being intentional under a description, the explication would be either circular or inadequate. For to the extent that it is true that our grip on the concept of action is sharpened by appeal to the concept of being intentional under a description in cases of negative actions, taking (I) as analyzing (or even as explicating) the concept of being intentional under a description is either question-begging or inadequate. If one takes the concept of action not to cover unintentional omissions, for instance, its use as an explicans of the concept of being intentional under a description is circular. For the concept of being intentional under a description is crucially involved in our having a firm enough grip on the concept of being an action not to apply to those cases. Alternatively, if one takes the concept of action to apply to cases of unintentional omissions, for instance (thus ignoring the way it is sharpened by the tie with the concept of being intentional under a description), the explication will fall short of capturing the concept of being intentional under a description because unintentional omissions are not intentional under any descriptions.

It thus seems that only two avenues are open if one wants to uphold (I). Either one abandons the reductive project and treats (I) non-reductively, or one treats (I) reductively but abandons some of the intuitions. For instance, if one were to uphold (I) as a reductive explication of the concept of action, one would have to give up the intuition that mistakes are actions, for instance.

4. Why abandon (I)?

An alternative avenue is to abandon (I) altogether. Why? We might first ask why one should adopt (I) in the first place (assuming the most plausible non-reductive explicatory reading). Since, as I argued, it would be implausible to construe (I) as representing even a non-reductive analytic link between the two concepts, the question that is reasonably asked is what it is that makes (I) even *prima facie* plausible (before one looks to cases where the concepts have a less clear application). The most plausible answer to this question is that (I) holds for cases that are sometimes considered to be

paradigms of actions: intentional movements.¹³ Thus, if one takes the category of intentional movements to be central to our concept of action, (I) will be immediately plausible. It will then seem worth the while to sharpen our intuitions on either one (as the reductivist would propose) or on both sides (as the non-reductivist would propose). The acceptability of (I) is thus conditional on the commitment to an investigation of the category of intentional behavior.

In the present project, I have undertaken an investigation of a related but in important ways distinct concept of *conduct*. The concept of conduct covers both intentional behavior as well as omissions (including unintentional omissions). Since one of the costs of upholding (I) involves tightening our concept of action to precisely exclude unintentional omissions, (I) is not even *prima facie* acceptable to someone who intends to understand the concept of conduct rather than intentional behavior.

A defender of (I) might object at this point that the initial plausibility of (I) ought to not only constitute reason for letting both categories involved be sharpened, but it also ought to throw doubt on the viability of the very project of trying to understand the concept of action as part of conduct rather than as part of intentional behavior. Such a theorist might argue that the intuitive appeal of (I) actually *shows* that unintentional omissions are not actions and so the very precept of the present project is called into question. In other words, it is illegitimate to reject (I) on the basis that it does not capture unintentional omissions, for (I) (fortified by its intuitive plausibility) actually demonstrates that unintentional omissions are not actions.

The objection fails. After all, the reason for thinking that unintentional omissions are not actions relies on their missing the connection with the concept of being intentional under a description. They have other conceptual connections that they share with what we recognize as actions (among them two prominent facts: we are held responsible for them and they form the basis on which we attribute character traits to people). It is not clear therefore that our unsharpened concept of action excludes the foundation of the

¹³ Recall from the Introduction that the category of intentional movements is extensional. It will thus be roughly coextensive with intentional actions and unintentional actions (which on Davidson's rendition are not extensional categories).

present project. What of the claim that we should reject the thesis that the concept of action covers unintentional omissions on the basis of the theoretical adequacy of (I)? I have already suggested that from the point of view of the present account, (I) does solidify *some* of our intuitions about the concept of action, viz. those that pertain to the idea of action as part of our intentional behavior. But there are other intuitions about the concept of action, viz. those that pertain to the idea of action as part of our conduct. It seems reasonable therefore to treat (I) as offering an adequate account of a narrower category than the concept of action that is under investigation here. But if so then it would be preposterous for however ardent a proponent of (I) to criticize the present account for not being narrow enough.

I have considered in some detail the thesis that to be an action is to be intentional under a description. We have seen that under the most plausible reading the thesis is understood as partly reporting but partly sharpening a conceptual connection that exists between the two concepts. I have suggested that there are cases (of negative actions) where the concept of being intentional under a description is used to sharpen our intuition of what negatively described performances ought to count as actions. It is on this ground that unintentional omissions are argued not to be our actions. And there are cases of spontaneous actions where the concept of action is used to sharpen our sense that there is some description under which they are intentional.

I have also suggested that what makes (I) so intuitively appealing is the fact that it unproblematically reports a preexisting conceptual connection for a range of cases, viz. intentional movements. Most intentional movements count as actions and they have some description under which they are intentional. For a theorist who aims to capture the concept of action understood as a unit of our intentional behavior, (I) will and ought to be a central thesis that is worth sharpening in any areas of unclarity. For a theorist who, as I do in this project, aims to capture the concept of action understood as a unit of our overall conduct, especially one of the costs of adopting (I) will be particularly unacceptable. Accepting (I) means that one has to deny that unintentional omissions and idle negative actions are to count as actions. But, as I insisted at the very outset, an account that attempts to capture the concept of action as a unit of conduct must include those cases. We are thus committed to rejecting (I).