

Chapter 4

What Do Deviant Causal Chains Deviate From?

Geert Keil

The problem of deviant causal chains is endemic to any theory of action that makes definitional or explanatory use of a causal connection between an agent's beliefs and pro-attitudes and his bodily movements. Other causal theories of intentional phenomena are similarly plagued.¹ The aim of this chapter is twofold. First, to defend Davidson's defeatism. In his treatment of deviant causal chains (DCCs), Davidson makes use of the clause "in the right way" to rule out causal waywardness, but he regards any attempt at specifying 'right' sorts of causal histories as hopeless and even harmful. To my mind, Davidson's defeatism contains a valuable insight, so I shall try to explain the reasons for it. Second, I shall try to answer a question that has often been ignored or passed over in the literature; namely the question of *what it is that DCCs deviate from*.

1. The problem of deviant causal chains

What is the problem of deviant causal chains? According to Davidson, the problem consists of the fact that "not just any causal connection between rationalizing attitudes and a wanted effect suffices to guarantee that producing the wanted effect was intentional. The causal chain must follow the right sort of route" (Davidson 1973, 78). Deviant cases have the status of counterexamples which threaten the adequacy of a causal account of the underlying intentional relation. DCCs are best considered as affecting the causal theorist's *definition* of intentional action. Davidson himself makes two attempts at such a definition. In his paper 'Agency', he says that a person performs an action "if and only if there is a description of what he did that makes true a sentence that says he did it intentionally" (Davidson 1971, 46). This explication is a pre-analytic one, since it appeals to the notion of doing something

¹ Deviant causal chains (DCCs) occur in causal theories of reference, knowledge and perception. A synoptic account of DCCs in the various theories of intentional phenomena is still missing. Most philosophers have concentrated on the case of action. The closest analogy is probably found in the causal theory of perception. Here I am thinking of cases of "veridical hallucination", as described by David Lewis (1986). In such cases the scene before my eyes happens to cause a hallucination of that very scene, such that the visual experience matches the scene before my eyes, and I still do not see that scene.

intentionally. This is, in Davidson's words, "to analyse the obscure by appeal to the more obscure – not as pointless a process as it is often thought to be, but still disappointing" (*ibid.*, 47). His second attempt, therefore, aims at finding "a mark of agency that does not use the concept of intention" (*ibid.*). We should keep in mind, however, that Davidson does not *dispute* the adequacy of equating "performing an action" with "doing something intentionally". What he seeks is a further analysis in terms of beliefs and desires. In "Intending", he arrives at the following analytic definition:

An action is performed with a certain intention if it is caused ...^[2] by attitudes and beliefs that rationalize it. (Davidson 1978, 87)

Note that Davidson's analysandum is not "performing an action" *tout court*, but performing an action "with a certain intention".

Now strictly, there can be no counterexamples to a definition. Definitions may be more or less useful, but they are neither true nor false. Things are different if the causal theorist's definition is meant to be an analysis of "doing something intentionally". Only by accepting this analysandum are DCCs counterexamples to the causal theorist's definition. If the causal theory of action does not find a way of ruling out deviant causal chains, it classifies a sequence consisting of a belief-desire pair, a bodily movement and a wanted effect as a case of intentional action which none of us would accept as being intentional. If the culprit sincerely declares "I didn't mean to do it like that", a theory of action should be capable of providing an explanation that accounts for this evidence. We are well advised to reject any definition or explanation that ties the agent down to an intention that he did not have. It seems desirable to rule out deviant cases by refining the conditions that have to be fulfilled for a wanted effect to be intentional.

Unfortunately only a few philosophers have turned this coin over to take a look at the other side. On closer examination, there are good reasons for *not* trying to exclude DCCs by refining the conditions. Davidson thinks that the price for excluding them would be too high, amounting to establishing strict (that is exceptionless) intentional laws. Davidson's well-known arguments for the anomalism of the mental will not be restated in this chapter. The arguments will however be touched upon, DCCs themselves being falsifying instances for such laws. I call the view that right sorts of causal histories cannot be specified *Davidson's defeatism*.

Calling DCCs "unavoidable" or "ineliminable", as many authors do, is ambiguous; hence one word of clarification. The issue is not whether nature is capable of thwarting our plans or not. Of course even our best intentions and plans can fail. The question under discussion is not whether DCCs can be avoided, but whether they can be ruled out by the causal theorist's analysis of "doing p intentionally", that is, whether the conditions the causal theory sets can be made strict enough to exclude all deviant cases, without excluding too much.

2 I have skipped the phrase "in the right way", which anticipates the problem of causal deviancy

2. Some examples

A person wants to commit suicide. She decides to play it safe. She swallows an overdose of sleeping pills, walks onto a river bridge, puts a rope around her neck, ties it to the handrail, loads her pistol and holds it to her forehead. Then she jumps down, inadvertently jerks the pistol, pulls the trigger, shoots through the rope, plunges into the icy water, is just saved from drowning, brings up the pills – and dies of pneumonia three days later.

Has she achieved her aim? Yes and no. The suicide had not intended to die in this way. Nevertheless, in this case one could perhaps reckon on some posthumous remark such as "So be it". Proceeding in such a circumspect way, the suicide has shown that her weariness with life clearly outweighed her preference for a particular way of dying. So, if her intention was to die no matter how, she has reached her goal. If her intention was to die in one of the four ways planned, she has not reached it.

Perhaps this case calls for special treatment. We might introduce a class of 'stoic' intentions characterized as follows: "to bring about the state p while accepting any conceivable incident, provided it helps to reach p". For such a class of intentions, the problem of DCCs would not arise. It is doubtful, however, whether such intentions deserve their name. The classical DCC examples are cases where the interfering event was unforeseen, perhaps even unforeseeable. Hence the event could not have been represented in my intention or plan. The retrospective "So be it"-style rationalization is a *mere* rationalization. It merely covers up the fact that what happened was not intended at all. For a *causal* theory of action, moreover, only those attitudes and intentions that an agent has formed beforehand are of use. The corresponding mental events are assumed to be the cause of the agent's bodily movement, and no later change in the agent's attitudes can contribute to that cause. What is done is done, and what was intended was intended.

The class of intentions that I was tempted to introduce in order to save the suicide (her rationalization, not her life) is presumably empty. We cannot allow such stoic intentions that embrace DCCs. Not bothering about unforeseeability would blur the distinction between things we do and things that happen to us.

Let us consider three familiar examples of DCCs.

- (i) *The sniper*: Daniel Bennett devised the case of an unpractised gunman who intends to shoot someone. His shot misses the victim, but startles a herd of wild pigs that trample the victim to death (cf. Davidson 1973, 78).
- (ii) *The driver*: A reckless car driver runs over a pedestrian who turns out to be the driver's uncle whom he wanted to kill (cf. Chisholm 1966, 37).
- (iii) *The mountaineer*: A mountaineer wants to get rid of the second man on the rope. This thought unnerves him and causes him to loosen his grip (cf. Davidson 1973, 79).

In all these cases a person's beliefs and desires cause and rationalize a bodily movement. It seems that all the conditions set by the causal theory of action are met, and yet we are not inclined to admit that the agent produced the effect intentionally.

The would-be sniper neither intended to startle the wild pigs, nor did he alert them intentionally; how could he if he did not even notice them? He wanted the victim's death, but he did not want *this* death. The pre-analytic judgement and the causal theory's analysis diverge.

3. Basic and nonbasic deviance

The cases hereto presented fall roughly into two groups. DCCs have been classified according to the *stage* in which the chain degenerates. In the wild pig case the causal chain degenerates outside of the agent's body, or, to put it in terms of temporal succession, after the agent has made his contribution. In the mountaineer case, the interfering event occurs inside of the agent's body. The example shows that short as well as long causal chains may degenerate. DCCs do not only occur in domino-style actions, or in extended, planned activities where several steps have to be taken. DCCs are undemanding; basic actions are affected as well.

Davidson speaks of "internal" and "external" DCCs. Bishop distinguishes "basic" from "non-basic" deviance, Mele "primary" from "secondary" deviance, Brand "antecedential" from "consequential waywardness". The distinctions are not completely equivalent. In particular, the internal/external distinction seems of limited interest, since meanderings that happen inside the body may be just as unintended and beyond one's control as external ones.³ I shall adopt Bishop's terminology: in cases of "basic deviance", the deviance "*affects the causal link between mental states and basic action*" (Bishop 1989, 133), while nonbasic deviance affects the causal link between the basic action and its further effects, as in the wild pig case. This distinction cuts across the agent's bodily limits, since a causal chain that degenerates inside an agent's body might do so *after* he has begun his basic action. Set out in this way, the distinction marks the difference between cases where, despite the DCC, *something was done intentionally*, whatever it was, and those cases where *no action was performed at all*. Both the sniper and the driver cases are clear cases of intentional action. Having startled the wild pigs unintentionally, the unpractised gunman can resort to a further, more basic description of his deed under which it was intentional, for example "shooting at the victim", or "pulling the trigger". Things are different with the uncomradely climber. Loosening one's grip as a result of nervousness is not intentional under any description, and thus cannot count as an action (according to Davidson's criterion, which is widely accepted in the debate).

Both cases, basic and non-basic deviance, are counterexamples to standard causal analyses of "doing p intentionally", that is performing an action *of a specified type*. If however the analysandum is "acting intentionally", instead of "doing p

³ To Stoecker, "it is hard to see why such a difference [that is, whether the causal chain goes wayward inside or outside the agent] should matter much. Imagine Bennett's would-be killer as a would-be doctor instead, who wants to self-cure his nausea with an injection of a substance that is in fact totally ineffective. Yet, clumsy as he is, he hurts himself so badly with the needle that the pain makes him vomit, which immediately cures his stomach. Although the causal chain leading to his recovery from nausea goes wayward inside of him, the example is obviously of the same type as the pigs case" (Stoecker 2003, 302).

intentionally", only the cases of basic deviance count as counterexamples (cf. Brand 1984, 18).

The most comprehensive classification of deviant chains was suggested by Christoph Lumer (see Lumer forthcoming, ch. 4.2). Lumer, incidentally, does not speak of "deviant causal chains". He makes it clear that strictly the term is a misnomer, since it does not capture cases where the deviation from an agent's intention does not rest in some causal link in a chain, but is due to a certain faulty assumption of the agent's. If, for example, the agent reaches his aim by mistaking someone for somebody else, what goes astray is not a causal process from his intentional state to his bodily movement or its further effects. To cover such cases of "tertiary waywardness" (Mele's term), Lumer suggests the collective term "deviant realizations of intentions". While the clarification is helpful, I shall stick to the inherited label "DCC" for the sake of convenience.

4. Some suggested solutions

4.1 Further specifications of the mental cause

When Davidson brought up the DCC problem, he was in search of an analysis of "doing something intentionally" that does not use the concept of intention. Davidson never did dispute that performing an action amounts to doing something intentionally, but until the mid-1970s he sought a *further* analysis that manages with beliefs and desires. The climber case may suggest that substituting intentions for belief/desire pairs will rule out basic DCCs. But this is not the case. Even if in Davidson's original example the climber's beliefs and desires immediately cause his nervousness, the story can easily be modified so that he loses control over his hands *after* he has formed a proximal intention. The question as to what it is that DCCs deviate from remains a hard one even if we extend our conceptual resources and allow ourselves to resort to intentions.

Several authors have proposed to specify in greater depth the mental cause of an action in order to rule out DCCs. They have taken the step from intentions to *plans*. As a matter of fact, many of our intentions are not merely directed towards the desired outcome of an action, but they include a detailed anticipation of how this state of affairs should be reached. These plans comprise, we might say, *means-directed* intentions, in addition to the *end-directed* intentions. So the suggestion is that something being brought about *in the right way* often means in the way *planned* and *anticipated*, so that DCCs would deviate from an agent's plans. In discussing Chisholm's driver example, Jerome Shaffer suggests the following treatment:

It is a case in which the end (the death of the uncle) requires a means. ... To deal with this, we must add the further condition that where it is necessary to employ some means in achieving an end, to bring about the end intentionally one must also bring about the means intentionally. (Shaffer 1968, 105)

This condition is not met in Chisholm's case, therefore, according to Shaffer, "the purported counterexample fails" (*ibid.*, 106).

True, achieving an end requires, in all cases of non-basic action, employing some means, and some ends involve taking more than one step. Taking the appropriate steps, however, may not always suffice to achieve one's aim. Over and above taking the appropriate steps, the agent must rely on the causal chains he has triggered running their usual course, that is, running as expected.

And if they do not? If something unforeseen happens to interfere, there are three possible cases to be distinguished:

- (a) The action fails, that is, the state of affairs that the agent wanted to bring about is not reached.
- (b) The agent manages to react to the incident. He succeeds in including the interfering event in his plan, or in making the requisite adjustments to the new situation in such a way that the desired result is attained nonetheless.

Neither (a) nor (b) are cases of DCCs. A deviant chain only develops if the agent cannot fit the incident into his action-plan, and if:

- (c) the desired state of affairs does materialize, but not in the way planned or foreseen.

In this last case the agent is merely, so to speak, nature's handyman. He works as a medium that nature uses to reach a state of affairs which, by pure chance, coincides with the state of affairs the agent tried to bring about in *his* way. Now, had we questioned Chisholm's driver about his intentions before the incident, his answer probably would have been "I want to kill my uncle." Well, there you are then! What else can we do but check whether the outcome of the action matches what the agent had sincerely declared to be his intention? Moreover, beyond what he sincerely *declared* to be his intention is there room for the question as to what he really intended?

Shaffer's suggestion is that DCCs deviate from an agent's *plans*, though not from his end-directed intentions. Insofar as action plans encompass a series of steps that have to be taken, and a reliance on foreseeable regularities, the way the agent caused the uncle's death obviously deviated from his plan. This deceptively simple answer is, however, not satisfactory for the following reason. On closer examination, every course of events deviates slightly from our plans and expectations, while not every such mismatch constitutes a DCC. To make things worse, our action plans do not anticipate the course of events down to the last detail, so that it is often not possible to cite a specific element in a plan that an unexpected incident deviates from. We shall return to this point in sections 6 and 7 below.

4.2 Gricean deference

Goldman, Armstrong and Mele employ a strategy which has been called "Gricean Deference", in allusion to a proposal that Paul Grice has made for the causal theory of perception.

Gricean Deference is an analytic technique. It is supposed to allow a philosopher to assert, say, ... that perception involves a particular causal relationship, while freeing him from any responsibility for an account of the cause's or the causal relationship's precise nature. For this account the philosopher *defers* at a certain point in his analysis to the specialist in the relevant special science, who is competent to speak on the causal relationship in question. (Cooper 1976, 91)

Applied to action theory, the proposal is to supplement the condition "if caused by beliefs and desires" with the clause "in the right way", or "in the normal way", while *deferring*, or *delegating*, any further specification of that normality clause to some special science. This is exactly what Goldman does:

[P]recisely what is this 'characteristic' mode of causation by which wants and beliefs cause intentional action? ... A complete explanation of how wants and beliefs lead to intentional action would require extensive neurophysiological information, and I do not think it is fair to demand of a *philosophical* analysis that it provide this information. (Goldman 1970, 62)⁴

The first step in this move is to add a normality clause. It is indeed tempting to say that DCCs deviate from the "usual" or "normal" way things go. But what does this mean? The mere words "normal", "usual", "right", "wrong", "deviant" or "wayward" are of no help here. Mother Nature draws no distinction between normal and deviant causal chains. Nature knows no right or wrong ways to bring things about. Whatever may go wrong with our actions, the course of events will not violate laws of nature. Various participants in the debate have observed that "the notion of deviance has not been given a sense in relation to causality" (Mitchell 1982, 353).⁵

A comparison with the difference between the desired and the undesired effects of a medicine might be illuminating. Nature does not recognize this difference either. The standard phrase to appear in the instructions reads: "Besides their main effect, medicines may have unwanted effects, so-called side effects." This phrase is very instructive. The difference between effects and side effects is exclusively that between wanted and unwanted effects. The property of being unwanted does not correspond to a distinct physical property. Side effects do not constitute a natural kind in any physical science.

But perhaps a DCC's being deviant does? Armstrong tries to reinforce Goldman's line of argument by using a computing analogy. His solution is very simple: "P does not bring about Q as a result of standard computing practices", due to a "disorganization of the computer's internal processes In the same way, we have within ourselves certain 'mechanisms' with certain powers", and these mechanisms may malfunction (Armstrong 1981, 84-5). While Armstrong employs the computer model of mind, Ginet (1990, 41) and Lumer simply invoke "reliable" or "sufficiently match-ensuring" mechanisms.

⁴ Surprisingly, Myles Brand imputes the view "that specifying causal normalcy is a scientific matter" (Brand 1984, 19) to Davidson. One wonders which passage in Davidson may have abetted this misreading. As far as I can see, Davidson has never encouraged scientific investigations into right sorts of causal histories.

⁵ Cf. Kenny 1975, 121; Searle 1983, 139; Brand 1984, 19; and Føllesdal 1985, 322.

But it is hard to see in which way the “mechanisms” of practical reasoning should have malfunctioned in the example Armstrong discusses (the climber case). Up to the output of the deliberation system everything works fine. Why should neurophysiologists sort out physiological processes that show no eccentricities whatsoever, except that they sometimes form links in causal chains that are a headache for some philosophers of action? And how *could* neurophysiologists sort such processes out? Our brains could not care less about DCCs.⁶

But the Gricean deferentist is not yet left empty-handed. Perhaps it is not the mechanisms of practical reasoning that malfunction, but some later physiological process. Mele suggests that in the climber case, there is an anomaly in the motor control system that can be identified physiologically. He devises a class of agents that are wired in such a way that they “can perform overt actions only in cases in which the acquisition of a proximal intention initiates the sending of motor signals to appropriate muscles”, so that “any alien intervention into normal processes of guidance simply shuts down the motor control system” (Mele 2003, 61 and 62). For such agents, the climber case would not threaten the causal theorist’s definition of intentional action, since the unnerved climber would either perform no action at all or at any rate not the action of dropping his partner.⁷ Mele admits that his solution is stipulative. He leaves it open whether the biological set-up of human agents is as he devises, that is, whether humans engage in overt action only if no alien intervention into normal guidance processes occurs.

It is noted above that the strategy of Gricean deference is prone to underestimate the problem of *identifying* a chain as deviant in the first place. It is one thing to recognize that Davidson’s climber has lost “control” or “guidance” of his movements, it is quite another thing to be able to *specify* this effect in a naturalistic way, that is, without invoking unanalysed agentive notions such as “guidance” and “control”. It should be clear that the causal theorist’s analysis “can be satisfactory only if it avoids any reference to actions, agent-causation, exercises of control, and the like” (Bishop 1989, 98).⁸

Now it is in the spirit of the Gricean deference strategy that the question of naturalistic specificability does not get definitively settled. According to Goldman (1970, 62), it is “not fair to demand of a *philosophical* analysis” that it provide the required neurophysiological information. But it seems that the task of specifying

6 For a more detailed criticism of Armstrong’s account, see Bishop 1989, 134–5.

7 Mele distinguishes various cases (2003, 59–60): (1) Either the agent’s intention does not result in motor signals being sent at all. (2) Or the motor signals do not reach his muscles. Instead they unnerve him and the relevant intention is quickly extinguished. (3) Or the intention is not extinguished, and in that case the agent’s nervousness temporarily deprives him of control over the motion of his hands. (4) The signals reach the muscles, his fingers start moving. This unnerves him, with the result that his grip loosens. In this case the movement of his fingers is not a “direct ballistic continuation” of his previous motions, hence no guidance, no control, no intentional action. In neither case, Mele claims, does the agent perform the action of dropping his partner.

8 This constraint is neglected in many answers to the DCC problem, for instance in Adams’s (1989) demand that the agent must not have the feedback control loop broken prior to his part in the action coming to a close.

a causal chain as deviant has two parts: first, classifying a sequence consisting of a belief-desire pair, a bodily movement and a wanted effect as unintentional, and second, trying to correlate this feature with some anomaly in the physiological or physical chain. It seems that the first part of the task can *only* be achieved by a philosophical analysis, and that the second question cannot even be addressed before the class of “deviant” cases has been specified in intentional vocabulary. And it might well turn out that any physiological predicate that covers the relevant malfunctionings would have to be wildly disjunctive.

If Mele is right, however, the deviant sequences do constitute a well-behaved physiological type. Since this matter cannot be settled here, I confine myself to the observation that *if* the naturalistic strategy of Gricean deference works, it does so only for cases of *basic* deviance, where the output of practical deliberation is prevented from causing the right kind of physiological process. It cannot work for tertiary waywardness, nor for nonbasic deviance, where all is fine both with the practical deliberation and with its causing a physiological and behavioural process that constitutes a full-blown action, as in the wild pigs case. Incidentally, many participants in the debate, including Davidson, think that basic deviance is harder to handle than non-basic deviance. I submit that exactly the reverse is the case. It is nonbasic deviance that defies a naturalistic solution.

4.3 The causal immediacy strategy

The last proposal I will scrutinize is the “causal immediacy strategy”, as John Bishop calls it (Bishop 1989, 138). The idea is that DCCs spread in a region which, in normal cases, must not exist, that is, in a causal gap between the mental antecedent and the bodily process caused. Myles Brand, one of the champions of the causal immediacy strategy, insists that there be no such gap:

An adequate Causal Theory must preclude the possibility of these types of interventions: there can be no causal space between the mental antecedent and the beginning of the physiological chain. (Brand 1984, 20)

As a remedy, Brand introduces the notion of “proximate causation” (*ibid.*). Roughly speaking, he simply disallows any intervening events. The only bodily movements that count as actions are those caused by immediately preceding mental events.

It is obvious that the causal immediacy strategy is designed specially for basic deviance. (Resorting to action plans, on the other hand, was designed for non-basic deviance.) The main problem with the causal immediacy strategy is that it is too restrictive. It excludes many clear cases of doing something intentionally. Normally, a spatiotemporal gap between mental antecedent and bodily movement does not preclude the latter being performed intentionally.⁹ Brand should be prepared to spell out *exactly what* must not happen in the meantime.¹⁰

9 See Keil 2001.

10 In addition, Brand’s account is dependent on a questionable notion of causal relata, as John Bishop draws attention to: “Have we ... any reason to suppose that [the very idea of

A variation on the causal immediacy strategy demands that the agent's intention accompany every stage of the action. John Searle demands that there be "continuous efficacy of intentional content":

We can now state the conditions necessary to amend the account as to eliminate all the deviant causal chains we have considered. A first condition is that there should be continuous efficacy of intentional content under its intentional aspects. (Searle 1983, 138)

It is not enough that one intended at some prior time, to perform the act; this intention must have persisted into the time when one acted, for the act to be intentional. (Mitchell 1982, 353)¹¹

Again, the proposal has to be evaluated twice, for basic and for non-basic deviance. If applied to the causal link between mental antecedents and basic action, the proposal coincides with the immediacy strategy just discussed. If applied to the subsequent stages, it becomes questionable what "continuous efficacy of intentional content" amounts to. How does the agent ensure such efficacy? Perhaps our rifleman, watching the wild pigs do their bloody deed, says to himself: "Oh, obviously the intentional content of my mental state isn't effective any more. Next time, I must equip it with more causal power!" The truth is that once the causal chain has left the agent's body, there is nothing whatsoever he can do to make his intention "continuously effective". The bullet flies simply too fast.

On a less absurd reading, Searle's demand is that the agent *keeps control* of the consequences of his basic action. Such a condition, however, is forbidden fruit for the causal theory of action, which tries to get by with event causality. Unanalysed talk of control violates this constraint (see Bishop's criticism quoted above).

Furthermore, Searle's condition is too strong. As shall be explained below, the *density* of actual courses of events precludes a perfect match with our intentions, anticipations or plans. And in many cases there is not even a fact of the matter about these incongruencies, since nature knows more than one way to fulfil our intentions. Hardly anything happens *exactly* according to our intentions and plans. This fact is not accounted for in any of the suggested remedies discussed hitherto. And it is this fact that makes our title question so hard to answer.

While the first two moves – further specification of the mental cause and Gricean deference – still operate within the conceptual framework of the causal theory of action, this is not so obvious for the causal immediacy strategy and for Searle's demand of continuous efficacy of intentional content. Since I am interested in the challenge that DCCs pose to the causal theory's original event-causal analysis of "doing something intentionally", I do not consider here any proposals that go beyond the Davidsonian metaphysics of singular causation between events. This is why I neglect teleological accounts, such as Schon's (1997), as well as the "differential explanation" strategy, as developed by Peacocke (1979) and Bishop (1989). Referring to "the existence of a suitable pattern of counterfactual dependence of

proximate causation] makes any more sense than a corresponding notion of 'proximate' points on a line?" (Bishop 1989, 139).

11. Thalberg, Audi and Frankfurt have made similar suggestions.

resulting behavior upon the content of the agent's basic intention" (Bishop 1989, 150) goes beyond the conceptual resources of an event-causal theory that views causality as an extensional relation. Besides, using the assessment as to whether, "had the agent's intention differed in content, the resulting behavior would have differed correspondingly" (*ibid.*), as a means of identifying DCCs, would require a superior kind of knowledge which is not available to the agent in his prior practical reasoning. The causal theory of action can invoke only those attitudes that cause the agent's movement. No later change in the agent's attitudes and no third party assessment can contribute to that cause. Perhaps the differential explanation strategy can even be harmonized with my own account (see below), but it cannot be moulded into a condition a Davidson-style causal theory could use to improve its original definition.

5. Defending Davidson's defeatism

At this point it might be helpful to recall the problem that Davidson responds to when bringing in, in addition to the agent's "primary reason", the causal relation between the reason and the action performed. The problem is that, in Davidson's words:

a person can have a reason for an action, and perform the action, and yet this reason not be the reason why he did it. Central to the relation between a reason and an action it explains is the idea that the agent performed the action *because* he had the reason. (Davidson 1963, 9)

The agent's rationalization must not, in other words, be a mere rationalization. The causal relation is brought into play in order to single out, among the reasons the agent had, the one that was effective in bringing about his behaviour. The appeal to causal efficacy is indispensable, since the rationalization by itself "provides no reason for saying that one suitable belief-desire pair rather than another (which may well also have been present in the agent) did the causing" (Davidson 1987, 42). The next thing we can expect the causal theorist to do is to *specify* this effective reason, or to tell us how it can be identified. But Davidson disappoints us here. He does not take any further steps to specify the effective reason. Among the reasons the agent had, he picks out the one which yields the most plausible explanation, but instead of specifying its position in the causal network of the physical world, he just *calls* it the cause, which is not particularly informative. The only principle Davidson employs here is the rule of thumb *best reason = strongest cause*. For an analytic definition of "doing p intentionally", he wishes to get by with the following elements: (1) the agent's bodily movement, (2) his primary reason (citing a pair of belief and pro-attitude), (3) the fact that the primary reason rationalizes the action, and (4) the fact that the primary reason (or rather, its occurrence) causes the action. We may complete the last condition by demanding that the cause operate "in the right way", but with this we run out of resources. The clause "all causal chains running normally" is nothing but a general *ceteris paribus* clause, serving to express our unwillingness to accept DCCs but giving no hint as to how to rule them out. More than these elements

we do not have, and, according to Davidson, we should not *want* to have. His reason is that improving on these conditions by further specifying them,

in a way that should eliminate wrong causal chains, would also eliminate the need to depend on the open appeal to causal relations. We would simply say, given these (specified) conditions, there always is an intentional action of a specified type. This would be understood as a causal law, of course, but it would not need to mention causality. Unavoidable mention of causality is a cloak for ignorance; we must appeal to the notion of cause when we lack detailed and accurate laws. (Davidson 1973, 80)

With this clarification, we have a most interesting interim result. Davidson has turned out to be anything but a prototypical causal theorist. Among the agent's rationalizing beliefs and pro-attitudes, he picks out a pair and calls it the cause of the action – and leaves it at that. A causal theory of action that makes *explanatory* use of the causal connection between the belief-desire pair and the action performed cannot leave it at that. Such a theory will have to *specify* the conditions required to eliminate the deviant cases. It will set out in search of a right sort of causal history argument. Davidson, however, is convinced that there cannot be such an argument. "What I despair of spelling out", he confesses, "is the way in which attitudes must cause actions if they are to rationalize the action". All we can say is that the effect has to be brought about "in the *right way*" (Davidson 1973, 79).

Davidson's despair in the face of DCCs is closely connected with his *anomalism* thesis of the mental. If we have to specify the conditions in such a way as to exclude DCCs, we would have to be in possession of something which does not exist, namely strict intentional or psycho-physical laws.

If the price for excluding DCCs by strengthening the conditions is too high, we are facing a dilemma, if not a paradox. A causal chain has badly degenerated, and we witness the agent's sincere declaration, "I didn't mean to do it like that." On the one hand, we are somehow convinced that this testimony has a truth-maker. On the other hand, given the agent's prior beliefs and desires, it is hard to explain why he is so dissatisfied. Did he not reach his declared goal? The Davidsonian background of the dilemma is this: on the one hand, various kinds of complications must be allowed so that human action can be the kind of thing it is – viz. intentional behaviour holistically embedded in a person's comprehensive pattern of interlocking beliefs and pro-attitudes, its interpretation constrained by normative considerations. Sequences of beliefs, pro-attitudes and bodily movements that fall under strict psycho-physical laws would no longer amount to the same thing, according to Davidson. Too many things can interfere with the anticipated courses of events to allow for a comprehensive list to be established. Due to the holism and to the normative character of the mental, there exist no true lawful generalizations of the form "given these (specified) conditions, there always is an intentional action of a specified type" (Davidson 1973, 80). If we found such a law, Davidson claims, we could be sure that we would have changed the subject.

On the other hand, a desired effect that was brought about in too bizarre a way ceases to be a case of "doing p intentionally". Now, the question of what it is that DCCs deviate from has to be distinguished carefully from the question of how the causal theorist's definition can be refined so that DCCs are excluded. Regarding the

latter question, I share Davidson's defeatism. It is a deplorable shortcoming of the DCC debate that the former question has often been skipped in favour of the latter.

6. The deep problem

The underlying problem, which makes the question of what DCCs deviate from so hard to answer, is the fact that hardly anything in the world happens exactly according to our intentions and plans.

The reason is not that the world is populated with evil demons who keep playing tricks on us. Quite the contrary: "Only a world that is regular in its sequence of happenings is a world in which intentions can be formed and executed. Only actions that have anticipated consequences are actions that can be performed for reasons" (Beck 1975, 116). It is not only that the *success* of our undertakings depends on Mother Nature's benevolence. Beck makes the stronger claim that in a chaotic world no beings could evolve that can *form* intentions in the first place.¹² However no exceptionless regularities or strict empirical laws are needed for that purpose. Defeasible regularities suffice.

The reason why the course of events typically deviates slightly from our expectations is that the world is *thicker*, or denser, than our mental representations of it. Our expectation of what will happen is, even in non-deviant cases, not detailed enough to be accurate. We do not anticipate our bodily movements right down to the last detail, let alone the way our movements will interact with the rest of the universe. Davidson is well aware of this fact:

[A]n intention cannot specify all the characteristics of the intended act that are relevant to its desirability. No matter how elaborately detailed an intention is, there are certain to be endless ways in which it could be realized that are unwanted and unintended by the agent. (Davidson 1985, 196)

Davidson speaks of intentions here, but it should be clear that with respect to foreseeability, plans are no better off than end-directed intentions. Now our plans and intentions cannot only be thwarted in endless ways, they can also be *fulfilled* by more than one course of events, for "[o]ne's intentions almost never specify an action so exactly that no action beside the particular action will satisfy them" (Morton 1975, 14).

The important distinction to be drawn here is this: our intentions and plans are directed towards actions of a certain *type*, falling under a certain *description*, while the *particular* action performed is underdetermined by the propositional content of our prior intentions.¹³ Action plans are not based on accurate anticipations of the future,

¹² Perry claims that even belief-desire psychology could not have evolved in a chaotic world. Human agents acting on beliefs and desires can rely on a certain benevolence on the part of Mother Nature, described by Perry as follows: "The actions we perform because we have certain desires and beliefs, are often of a sort that will promote the satisfaction of the desires if the beliefs are true" (Perry 1986, 194).

¹³ Cf. Davidson 1978, 96–7.

but on reasonable *abstractions*, which are made according to presumed relevance. In most cases we get by with these abstractions. Since we have not anticipated events in detail, we cannot compare them afterwards with our expectations point by point. Therefore it is, strictly, not correct to say that such details *deviate* from any element of our plan. If we did so, there would be no non-deviating chains left. As long as the countless unplanned details of an action performed are insignificant enough, there is simply no fact of the matter about “deviations”. Though the unplanned details *exceed* our plans, there is no specific element in the plan that they *deviate* from.

Now, in our sample cases of DCCs the unplanned details that exceed our plans can no longer be ignored. But again, it is, strictly, not correct to say that the disturbance concerns a particular element of the agent’s intention or plan. The agent has, after all, simply not considered this complication before. He did not even notice the wild pigs, therefore he could not form an intention to avoid the incident.¹⁴ Planning his action and anticipating the future happenings, he did rely on some unspecific expectation that everything would go smoothly, but he had not drawn up a list of possible complications. In that sense all action plans include an implicit *ceteris paribus* clause. Since the causal route by which the result is reached might always meander in an unforeseeable way, the agent is not able to name a desired end state such that, having reached it, he would, whatever the case, have to be content with what has happened, and to admit to have brought it about intentionally.

7. Spots of indeterminacy

Sometimes Mother Nature thwarts our plans, and some of these complications qualify as DCCs. But which ones? My leading question was not whether science can rid us of DCCs, but rather, what it is that DCCs deviate from. The fact that human agents are not Laplacean demons makes DCCs ineliminable, but it does not make them deviant.

Hitherto our enquiries have not been very encouraging. It seems as if we have been collecting answers to the question as to what DCCs do *not* deviate from. We have learned that they do not deviate from nature’s way of bringing things about, since nature knows no right or wrong ways to bring things about. I have maintained that, strictly, DCCs do not even deviate from an agent’s prior intentions and plans. We can perhaps say that they deviate from certain assumptions of normalcy *implicit* in the agent’s intentions or plans, but as long as the notion of a normal course of events defies analysis, this answer provides little illumination. It is typical of DCCs that we cannot cite a particular element of an agent’s intention that an interfering event clashes with.

The deep problem about DCCs is that nature knows of more than one way to meet our expectations, so that in many cases events *exceed* our intentions and plans without *coming into collision* with them. This fact makes our title question so hard to answer.

¹⁴ Stoecker makes the same point for cases of basic deviance: “The climber had no plans or expectations about the origin of his loosening the grip and consequently no false ones” (Stoecker 2003, 301).

Now literary theory has coined the illuminating notion of “spots of indeterminacy”. All fictional texts contain such indeterminate spots, or blanks, that is zones where there is no telling, on the basis of what the text says, whether a certain object or situation has a certain property or not.¹⁵ Take Emma Bovary. Does she have a mole on her right shoulder or not? If Flaubert does not tell us, nobody will. If the novel does not contain the information, the question must remain open. There is simply no fact of the matter. The real world, on the other hand, contains no such blanks. Every one of us either has a mole on the right shoulder or not (or more than one). Unlike fiction, the real world is fully determinate.

Now our plans and intentions are like fictional texts in that they contain indeterminate zones. The content of a future-directed mental representation is not as dense as the real world. The representation contains blanks which the real world can fill in different ways, some of which constitute DCCs, some of which do not. But in both cases the filled-in details do not clash with any particular element of the intentional content of the agent’s prior attitude. The agent’s intentional attitude simply leaves the details open, just like Flaubert’s novel leaves open the exact constitution of Emma Bovary’s shoulder.

8. What deviant causal chains deviate from – and why this is bad news for the causal theory of action

Having touched upon many related issues, we have still found no convincing reply to our title question. On the one hand, the climber, the driver and the rifleman have reasons to be somewhat dissatisfied with the course of events. On the other hand, every attempt has failed to specify conditions for a ‘normal’ course of events from which the eccentric killings were deviations. And worse still: as Davidson sees it, specifying conditions for a normal course of events that deviant chains deviate from would amount to the establishment of a strict intentional law of the sort “Given these (specified) conditions, there always is an intentional action of a specified type” (Davidson 1973, 80), and this venture is hopeless, according to anomalous monism.

The question remains of how tight this package is tied up. The impression is hardly avoidable that there is a point to the agent’s sincere assertion that he did not mean to do it like that, and it is only fair to demand of a philosophical analysis that it accounts for this assertion. Otherwise we lack any reason to call the cases deviant. It seems desirable, though, to have an analysis which does not require a strict conceptual or nomic link.

I dare say that the solution has been staring us right in the face all the time. It is worth noticing that the agent’s autobiographical report we wish to account for is always given in the past or perfect tense: “I *didn’t* mean to do it like that!” It is true: the rifleman did not kill his victim intentionally in this roundabout way, though nothing in his prior intention explicitly ruled out this course of events. When

¹⁵ “We find such a place of indeterminacy wherever it is impossible, on the basis of the sentences in the work, to say whether a certain object or objective situation has a certain attribute” (Ingarden 1973, 50).

starting his action, the agent was not in a position to formulate an intention detailed enough to exclude every possible complication. *In retrospect*, he can tell whether *this particular* course of events was intentional or not. So I wish to propose the following answer:

DCCs deviate from *that counterfactual course of events which would have corresponded to the retrospectively specified formulation of the agent's intention*.

Beforehand, the agent could take into consideration potential incidents only in the form of a general *ceteris paribus* clause. When the action is completed, his epistemic position is better. Now he knows (though not necessarily!) which of the vast number of potential complications has occurred, and he can substitute an explicit specification of his intention for the general *ceteris paribus* clause. For the prior reservation "The course of my action will be intentional if the causal chain unravels normally", he can now substitute: "The course of my action was unintentional because the wild pigs intruded."

What DCCs deviate from can only be told *ex post actu*. This result accords with the insight, which is not alien to Davidson (cf. 1963, 16), that explanation of action contains an irreplaceable *ex post*-element. We should not hesitate to apply this insight to the problem of DCCs. We cannot tell what they deviate from without referring to the agent's retrospective specification of his intention, which he was not in a position to give beforehand.

I insist on the legitimacy to speak of a *specified* intention instead of an *altered* or a *revised* one. Surprised by the wild pigs' interference, the agent narrows down his prior intention to kill the victim, but he does not withdraw it. If he did so, he would be lying about his past intentional states.

In speaking of a retrospective specification, I do not want to inaugurate a distinction between prior intentions and *ex post actu* intentions. There cannot be such things as *ex post* intentions, as long as we regard intentions as mental causes of actions, as the causal theory of action does. What is done is done, and what was intended was intended. The past cannot be undone, nor can the causal efficacy of past events. The distinction to be drawn is that between *intending to do p and doing it* on the one hand and *doing p intentionally* on the other. The so-called "simple theory" equates both phenomena, but the DCC cases reveal that this cannot be right. The adverb "intentionally" qualifies the agent's attitude towards a given *particular* action, whereas the mental act of intending, as Davidson says, "cannot single out a particular action in an intelligible sense, since it is directed to the future" (Davidson 1978, 99). The *ex post* judgment, I wish to add, is *de re* and *indexical*, while the content of the prior intention is *de dicto* and *descriptive*. And only because nature can find a way to place its complications into the *unspecified* parts of our plans and intentions, it is non-paradoxically true that not every action that arose from a prior intention has been done intentionally.

This is bad news, though, for the *causal* theory of action. The causal theorist cannot gain from the suggested account. In need of reputable causal relata, he has restricted himself to prior mental episodes. The agent's attitudes, or rather changes

in his attitudes,¹⁶ function as the cause of his bodily movement, and no subsequent modification can contribute to that cause.

For a Davidson-style causal theory, the agent's prior attitudes perform a double duty: they play their role as causes, and they are supposed to *explain* the action. However, the explanatory force of the assertion that the agent acted exactly for that reason which caused his action was modest enough, since Davidson's causal theory lacks any independent characterization of this cause. Davidson counters this objection with the claim that rationalization itself "is a species of causal explanation" (1963, 3), that is, that citing the appropriate belief-desire pair is *per se* causally explanatory. In view of our treatment of DCCs, however, it becomes hard to see in which way the *ex post* rationalization should pass for a causal explanation. Causal and rational explanations fall apart again. In deviant cases the belief-desire pair was still the cause of the action, and in cases of nonbasic deviance the beliefs and desires even *rationalize* the action under a certain description (for example "shooting at the victim"), but they are too unspecific to rule out DCCs. Nevertheless, we are not left empty-handed. We are capable of sifting out a particular causal chain as deviant. Our account is based on an *ex post*-specification of the unspecified "right way"-clause. It is only the causal theorist who is left empty-handed. Davidson can stick to his view that "the propositional attitudes are by nature explanatory" (1986, 206) only if he parts company with the orthodox causal theory of action, which must comply with the agent's prior attitudes. What enables us to sift out certain causal chains as deviant are the agent's attitudes, but not the ones he had when he began his action.¹⁷

9. Three objections

I would like to conclude with an attempt to counter three objections that may be raised against my proposal.

¹⁶ "[I]t is *changes* in the attitudes, which are events, which are the often unmentioned causes. ... [T]he cause of the action was the *advent* of one or both of the belief-desire pair" (Davidson 1993, 288).

¹⁷ I have covered up a point of disagreement. When he introduces the DCC problem, Davidson says that "not just any causal connection between rationalizing attitudes[!] and a wanted effect suffices to guarantee that the wanted effect was intentional" (1973, 78), while in other places he says that the beliefs and desires *would have rationalized* the action *if* they had caused it in the right way (1973, 79). In the latter, 'hypotactic' account he makes the rationalization *dependent* upon the existence of the right causal chain (rationalized *if* caused in the right way), while in the former, 'paratactic' account the rationality condition and the causality condition are being evaluated separately (rationalized *and* caused in the right way). Though Davidson does not seem to have noticed the tension between both formulations, the hypotactic account plausibly reflects his real opinion (for the details, see Keil 2002). But the hypotactic account is untenable. The rifleman's practical deliberation was unobjectionable, and what he *did*, that is, aiming and pulling the trigger, was perfectly rational in the light of his beliefs and desires. The fact that the causal chain meandered does not affect the rationalization. His attitudes do rationally explain his basic action, even if they do not make it the case that he aroused the wild pigs intentionally.

(i) First, how can I account for *hidden* DCCs? We can well imagine DCCs which are only detected after years, or never, or not by the agent himself. Let us assume that our gunslinger has been successful. No wild pigs. The victim is hit right in the head. Many years later it turns out that the shot was a ricochet, which had bounced off a nearby rock. We might ask what this causal chain deviates from until the detour is discovered. And what if it is never discovered? Must we conclude that there is no deviancy in such cases?

Well, I did not claim that DCCs are *necessarily* detected afterwards. I claimed that they *can only* be detected afterwards. Some DCCs may be hard to trace, but this has no bearing upon the asymmetry between what the agent can tell beforehand and what he can tell afterwards. Now imagine that the waywardness is detected by some observer, whereas the agent remains ignorant. In such a case the observer knows better what actually happened. This knowledge alone, however, does not put him in a position to judge that a DCC has occurred. Identifying a DCC is not merely a matter of discovering unknown physical facts. Ascertaining which meanderings count as DCCs requires a re-evaluation of the agent's attitudes in view of the incident. Cases of insincerity and self-deception set aside, this re-evaluation must not contradict the agent's own judgement.

Such an investigation, however, concerns the verification of a DCC, not its definition. Rejecting verificationism, as we should do, we must not infer from contingent problems of verification that there is no fact of the matter about the phenomenon to be verified.

(ii) The second objection stems from a further counterexample, which is another variation on the wild pig case. It goes like this: what if the rifleman had formulated his intention in a different way? What if he had declared beforehand that his intention was to kill the victim with a well-aimed shot to the head? Furthermore, let us assume that he had noticed the wild pigs, and had even considered the threatening complication, but then discarded it as too unlikely. He shoots, his assumption proves wrong, the wild pigs do their bloody deed. Would not, contrary to my claim, the course of events in this case clearly deviate from a specific element in the agent's prior intention? And would not the agent have been in a position to tell in advance that, if the wild pigs should intervene, his intention would not be fulfilled?¹⁸

Both questions must be answered in the affirmative. But I am not convinced. The objection is beside the point, since the present case would no longer be a DCC. For something to be a DCC, the state of affairs brought about must fulfil an agent's prior intention. This condition is met in all the examples I used, but not in the present case. Here we are confronted not with a DCC, but with a case where the state of affairs that the agent intended to bring about is not reached (see above, section 4, case (a)). If the agent had considered the threatening complication beforehand, he would have intended either to include it in his plan or to avoid it. In both cases, the complication would not have constituted a DCC. It goes without saying that in these cases, as always, another complication might take place that bypasses even the more specific intention.

¹⁸ This objection was raised by Andreas Kemmerling. Discussing it with him helped me to clarify my position.

(iii) Third, one may raise doubts as to whether the agent's retrospective re-evaluation of his intention is grounded in fact. If nothing in his prior mental state ruled out the wayward chain, then which good reasons can the agent cite for the claim that he "didn't mean to do it like that"? Does he enjoy a kind of first person authority when he specifies his intention in retrospect? Are there any objective constraints, or can he say just anything? Perhaps he has just changed his mind?

Let us assume that the objection does not question the credibility of autobiographical reports *in general*. True, an agent who prefers to dissociate himself from his deed when seeing his victim in a pool of blood may be *insincere*. But such cases are of no interest for the present debate. If the agent has incorporated the wild pig complication into his secret plan, the complication would no longer constitute a DCC. We are dealing with cases instead where we have every reason to buy the agent's tale.

One fact is beyond dispute, *viz.* the agent's prior intentional state. The past cannot be undone, and if the content of the agent's intention was, say, "I'm going to kill this guy", then no retrospective judgment can be sincere that renounces this intention. This is why I speak of a retrospectively specified intention instead of a revised one.

Given my story about blanks in our anticipations, that is, given that future-directed intentions can be fulfilled by a certain *range* of courses of events, the question obviously matters whether the deviation was a *considerable* deviation. In the literature, the case is discussed of a soccer forward who kicks the ball hard towards the opposing goal with the intention to score. The goalkeeper, however, manages to touch the ball,

thus deflecting it slightly from a straight path, but not nearly enough to cause it to miss the goal. It does not seem that *this* unexpected deflection makes it wrong to say that *S* kicked the ball into the goal intentionally. (Ginet 1990, 79)

I agree. If the player's intention to score the goal is fulfilled, then the exact trajectory of the ball is a matter of indifference to him. But this holds only within certain limits. What are those limits? What kind of mismatch is so significant that it makes the result unintentional? Ginet simply demands that "one must not be too lucky" (*ibid.*, 78). Lumer argues as follows: it is generally permissible to make probabilistic assumptions about one's prospects of success. Humans act in a world where success is never guaranteed, hence intentional action must not be incompatible with probabilistic calculations. But some complications are beyond that statistical range. Lumer offers the following criterion: would the agent, had he known about the threatening complication before, have proceeded in the same way, or would he have changed his plan and tried in another way? The soccer player would arguably not have proceeded differently, hence he scored the goal intentionally (see Lumer forthcoming, 222).

I feel free to adopt this idea, since it is in the spirit of my own proposal anyway. None of us knows the future. The piece of counterfactual reasoning that Lumer makes use of – *Would knowledge of the incident have made a difference?* – is only available *ex post actu*. Beforehand, in his practical deliberation, the agent is simply not in a position to judge whether knowledge of *this* deviation would affect his action plan

or not. And if the agent had considered the threatening complication before, it would not have constituted a DCC (see my reply to the second objection).

His foresight being limited, the agent can assess the particular course of events only *ex post actu*, and decide *of it*, hence *de re*, whether he brought it about intentionally or not. But he *might* have said in advance which course he approves of, if only a Laplacean demon had drawn his attention to the threatening complication. And this is why the piece of counterfactual reasoning Lumer offers resolves the doubt as to whether the agent's retrospective specification of his prior intention is grounded in fact. Though there is, strictly, no fact of the matter about whether his prior attitudes ruled out the deviancy, there is a *fact of the mind*, as it were. The gap between both facts can be closed by applying the principle of the supervenience of the mental: a world in which the assertion "I didn't mean to do it like that" is insincere and hence false, must contain a physical difference as well, wherever it is located. In a somewhat ethereal sense of "fact of the matter", we may even say that there is such a fact *ante actum*, namely in the agent's *dispositions*. Nobody asked him then, but this does not mean that the answer would have been arbitrary. I submit that the mental state he was in when deliberating his action fixed which answer he *would have given*, had a being endowed with foreknowledge questioned him about the complication. This prior disposition serves as a constraint for his retrospective re-assessment of his intention.

References

- Adams, Frederick (1989), 'Tertiary Waywardness Tamed', *Critica*, 21, 117–25.
- Armstrong, David M. (1981), 'Acting and Trying', in David M. Armstrong, *The Nature of Mind*, Brighton, Sussex: Oxford University Press, pp. 68–88.
- Beck, Lewis White (1975), *The Actor and the Spectator*, New Haven, CT and London: Yale University Press.
- Bishop, John (1989), *Natural Agency: An Essay on the Causal Theory of Action*, Cambridge: Cambridge University Press.
- Brand, Myles (1984), *Intending and Acting: Toward a Naturalized Action Theory*, Cambridge, MA: Cambridge University Press.
- Chisholm, Roderick M. (1966), 'Freedom and Action', in K. Lehrer (ed.), *Freedom and Determinism*, New York: Random House, pp. 11–44.
- Cooper, W.E. (1976), 'Gricean Deference', *Metaphilosophy*, 7, 91–101.
- Davidson, Donald (1963), 'Actions, Reasons, and Causes', in Davidson 1980, pp. 3–19.
- (1967), 'Causal Relations', in Davidson 1980, pp. 149–62.
- (1970), 'Mental Events', in Davidson 1980, pp. 207–25.
- (1971), 'Agency', in Davidson 1980, pp. 43–62.
- (1973), 'Freedom to Act', in Davidson 1980, pp. 63–81.
- (1974), 'Psychology as Philosophy', in Davidson 1980, pp. 229–44.
- (1978), 'Intending', in Davidson 1980, pp. 83–102.
- (1980), *Essays on Actions and Events*, Oxford: Oxford University Press.
- (1985), 'Replies', in Bruce Vermazen and Merrill Hintikka (eds), *Essays on Davidson – Actions and Events*, Oxford: Oxford University Press, pp. 195–229, 242–52.

- (1986), 'Judging Interpersonal Interests', in J. Elster and A. Hylland (eds), *Foundations of Social Choice Theory*, Cambridge, MA: Cambridge University Press, pp. 195–211.
- (1987), 'Problems in the Explanation of Action', in P. Pettit, R. Sylvan and J. Norman (eds), *Metaphysics and Morality*, Oxford: Oxford University Press, pp. 35–49.
- (1990), 'Representation and Interpretation', in K.A.M. Said et al. (eds), *Modelling the Mind*, Oxford: Oxford University Press, pp. 13–26.
- (1993), 'Replies', in R. Stoecker (ed.), *Reflecting Davidson: Donald Davidson Responding to an International Forum of Philosophers*, Berlin and New York: de Gruyter and Hawthorne.
- (1995), 'Laws and Cause', *Dialectica*, 49, 263–79.
- Føllesdal, Dagfinn (1985), 'Causation and Explanation: A Problem in Davidson's View on Action and Mind', in E. LePore and B. McLaughlin (eds), *Actions and Events: Perspectives on the Philosophy of Donald Davidson*, Oxford and New York: Blackwell, pp. 311–23.
- Ginet, Carl (1990), *On Action*, Cambridge: Cambridge University Press.
- Goldman, Alvin I. (1970), *A Theory of Human Action*, Princeton, NJ: Princeton University Press.
- Ingarden, Roman (1973), *The Cognition of the Literary Work of Art*, Evanston, IL: Northwestern University Press.
- Keil, Geert (2000), *Handeln und Verursachen*, Frankfurt am Main: Klostermann.
- (2001), 'How Do We Ever Get Up? On the Proximate Causation of Actions and Events', *Grazer Philosophische Studien*, 61, 43–62.
- (2002), 'Über eine Inkonsistenz in Davidsons Handlungsdefinition', in A. Beckermann and C. Nimtz (eds), *Argument und Analyse – Sektionsvorträge*, Paderborn: Mentis, pp. 668–78, <http://gap-im-netz.de/gap4Konf/Proceedings4/proc_txt.htm>.
- Kenny, Anthony (1975), *Will, Freedom and Power*, Oxford: Oxford University Press.
- LePore, Ernest and Brian P. McLaughlin (eds) (1985), *Actions and Events: Perspectives on the Philosophy of Donald Davidson*, Oxford and New York: Blackwell.
- Lewis, David (1986), 'Veridical Hallucination and Prosthetic Vision', in David Lewis, *Philosophical Papers*, vol. II, New York and Oxford: Oxford University Press, pp. 273–86.
- Lumer, Christoph (forthcoming), *Kognitive Handlungstheorie*, ch. 4.2: 'Abwegige Absichtsrealisierung'.
- Mele, Alfred R. (1987), 'Intentional Actions and Wayward Causal Chains: The Problem of Tertiary Waywardness', *Philosophical Studies*, 51, 55–60.
- (2003), *Motivation and Agency*, Oxford: Oxford University Press.
- Mitchell, Dorothy (1982), 'Deviant Causal Chains', *American Philosophical Quarterly*, 19, 351–3.
- Morton, Adam (1975), 'Because He Thought He Had Insulted Him', *Journal of Philosophy*, 72, 5–15.

- Peacocke, Christopher (1979), *Holistic Explanation*, Oxford: Oxford University Press.
- Perry, John (1986), 'Circumstantial Attitudes and Benevolent Cognition', in John Perry, *The Problem of the Essential Indexical and Other Essays*, New York and Oxford: Oxford University Press, pp. 193–204.
- Searle, John R. (1983), *Intentionality: An Essay in the Philosophy of Mind*, Cambridge, MA: Cambridge University Press.
- Sehon, Scott R. (1997), 'Deviant Causal Chains and the Irreducibility of Teleological Explanation', *Pacific Philosophical Quarterly*, 78, 195–213.
- Shaffer, Jerome A. (1968), *Philosophy of Mind*, Englewood Cliffs, NJ: Prentice Hall.
- Stoecker, Ralf (ed.) (1993), *Reflecting Davidson: Donald Davidson Responding to an International Forum of Philosophers*, Berlin and New York: de Gruyter and Hawthorne.
- (2003), 'Climbers, Pigs and Wiggled Ears', in S. Walter and D. Heckmann (eds), *Physicalism and Mental Causation*, Exeter and Charlottesville, VA: Imprint Academic, pp. 295–322.
- Vermazen, Bruce and Merrill B. Hintikka (1985), *Essays on Davidson – Actions and Events*, Oxford: Oxford University Press.