

Ezio Di Nucci

Automatic Actions: Challenging Causalism

Abstract:

I argue that so-called *automatic actions*—routine performances that we successfully and effortlessly complete without thinking such as turning a door handle, downshifting to 4th gear, or lighting up a cigarette—pose a challenge to causalism, because they do not appear to be preceded by the psychological states which, according to the causal theory of action, are necessary for intentional action. I argue that causalism cannot prove that agents are simply unaware of the relevant psychological states when they act automatically, because these content-specific psychological states aren't always necessary to make coherent rational sense of the agent's behaviour. I then dispute other possible grounds for the attribution of these psychological states, such as agents' own self-attributions. In the final section I introduce an alternative to causalism, building on Frankfurt's concept of *guidance*.

1. Introduction

“[. . .] und schreib getrost ‘Im Anfang war die Tat’” (Wittgenstein, *On Certainty*, §402)¹

We think less than we think. An awful lot of what we do either is automatic or it involves automatic performances and processes. Think of what you have done so far today: getting out of bed, going to the toilet, putting the kettle on, turning on the radio, brushing your teeth, getting dressed, walking to your office. These, in turn, involved a lot of turning handles, taking steps, raising arms, pushing buttons.

What all those sets of movements have in common is what one could call *mindlessness*: you did not think about these movements, nor did you need to. Mindlessness then distinguishes these movements from others: finally confessing to a wrong-doing; holding on to the rope from which your best friend is hanging; driving through a snow-storm on a mountain road, at night. Those actions are not automatic: they require a lot of thinking, wondering, pondering, deliberating, and hesitating; much attention, care, controlling, making sure. They

¹ Anscombe and von Wright leave the passage in German in the English translation (by Anscombe and Paul) of *On Certainty* (1969), translating it only in a footnote: “[. . .] and write with confidence ‘In the beginning was the deed’”. “Im Anfang war die Tat” is from Goethe's Faust.

require both mental and physical effort and strain. Also, those actions, differently from many of our automatic performances, will not be easily forgotten.

There are performances we can effortlessly and successfully complete without paying attention to or becoming aware of them: turning a door handle, *skills* like downshifting to 4th gear, or *habits* like lighting up a cigarette. When we act automatically we don't, consciously or unconsciously, think; nor do we need to think. Through practice, we become confident enough with our automatic performances that we can spare much of the cognitive resources which are normally required by novel or unfamiliar activities.

When we act automatically, we don't deliberate in advance over whether to act in that way. Nevertheless, we act deliberately. When we act automatically we don't first formulate in our mind a goal that our action is supposed to achieve. Nevertheless, what we do automatically is often goal-directed. When we act automatically we don't reason to decide what to do. Nevertheless, we normally act both rationally and reasonably. When we act automatically, we don't first form prior intentions to so act. Still, what we do, we normally do intentionally.

While psychologists have always been very interested in automaticity (from James's *Principles of Psychology* (1890) to Bargh's very influential work over the last two decades (1996; 1999; 2006; 2008)), philosophers have seldom paid it much attention.² One notable exception is Ryle: "When we describe someone by doing something by pure or blind habit, we mean that he does it *automatically* and without having to mind what he is doing. He does not exercise care, vigilance, or criticism. After the toddling-age we walk on pavements without minding our steps." (Ryle 1949, 42—emphasis mine)³

Not every automatic movement of ours is something we do, or an action. The kinds of automatic actions described above appear to differ from other automatic movements: reflexes like eye-blinking; tics; nervous reactions like sweating; biological processes like digestion and heart-beat; bodily changes like hair-growth; unconscious movements like sleep-walking; O'Shaughnessy's (1980, Ch. 10) "sub-intentional acts", such as the movements of one's tongue. The difference between these kinds of movements and the automatic performances described above seems to be *agency*. Even though we have some degree of indirect control over our heart-beat or hair-growth, those aren't things we do. While taking a sip of beer or typing the word 'beer' might also go unnoticed, but it is us who *do* those things—they don't merely *happen to us*.

In this paper I argue that *automatic* actions pose a problem for causal theories of agency (from Davidson's *Actions, Reasons, and Causes* downwards) because they cannot be reconciled with the 'mental' causal story that those theories tell. The intellectualist picture that founds causal theories is questioned by automatic actions so that causalism cannot be stretched to account for the obviously intentional and rational character of these performances: automatic actions do not fit the causal view because they challenge the required attribu-

² For my critical discussion of Bargh's work, see Di Nucci 2008 and forthcoming.

³ See also Dreyfus 1984; 1988; 2005, who focuses particularly on skill, and more recently Rietveld 2008.

tion of content-specific psychological states as causes in every case of intentional and rational agency; or so I shall argue.

2. Details? Aristotle Didn't Think So

Who cares, you might think, about turning handles, taking steps, and pushing buttons? Even if the causal view failed to account for *these* sorts of actions, these are certainly not the kind of deeds that *define* agents; and therefore we shouldn't trouble ourselves with whether our theory can deal with such details.

Two points here: firstly, automatic activities feature in most, if not all, deeds. If one cared to understand actions in terms of components or parts (Thalberg 1977; Thomson 1977; Ginet 1990), one could easily see how most if not all actions have some automatic component or aspect. And if, on the other hand, one accepts—as I do—the prevailing picture of action individuation (Anscombe 1957; Davidson 1971; Hornsby 1980) according to which we ought to speak of movements in terms of their descriptions in such a way that more than one description can capture the same set of movements so that multiple action-descriptions can refer to the same action, one will find that most actions are captured by automatic descriptions: so that when Oswald killed Kennedy, he might have well pulled the trigger automatically.

But there is a more poignant *Aristotelian* point to make about the importance of automatic actions. In the *Nicomachean Ethics* (Book II) Aristotle's idea is that the virtuous person is the one who *naturally* opts for the good deed; the one who doesn't have to decide or deliberate over which is the good deed. The virtuous deed is, in short, the one that the agent does not need to think about: it is only when virtue becomes second nature that the agent becomes virtuous: "his action must proceed from a firm and unchangeable character" (*Nicomachean Ethics*, Book II, 4—Ross's translation). The agent, in a slogan, can't *choose* virtue: she must *be* virtuous, as the result of having been habituated to virtue in her upbringing (see McDowell's concept of *Bildung* (1996, 84)).

The agent might perform an act that is in accordance with virtue, but if that act does not spring directly out of the agent's "unchangeable character" (her second nature), then her action won't be virtuous. The intuition behind this is, I think, that actions that the agent performs naturally, effortlessly, without hesitation, spontaneously, are *truer* to the agent's self—and many of those actions will probably be automatic ones (that is, they will be captured by many automatic descriptions). Only if the agent's adherence to virtue is true, spontaneous, and genuine can her actions be virtuous and the agent virtuous. Otherwise, according to Aristotle, the agent is merely continent.

This is what it means for automatic actions to be *truer* to the agent's self: they don't tell us who the agent *aspires* or wishes to be; they don't tell us what the agent's *ideal* self is (see Smith 1996). They tell us who the agent *actually* is; who she has become through the years; whom she has made herself into. The risk of looking at agency only through the lenses of motivation is that we get

a picture of who the agent wants to become rather than of who she is or has become.

Since automatic actions spring from the agent without the medium of thought, then it is only natural to suppose that they are more *the agent's own* actions than those that have been thought through. The less does an agent think about ϕ -ing, the more is ϕ -ing the agent's own: a *truer* expression of who the agent is, because it is one which wasn't mediated, nor needed to be mediated, by thought.⁴

To understand this it helps to go back to the process of *Bildung*: it is a process of internalisation; it is a process of appropriating particular performances that the agent has grown comfortable, and confident, with. Those performances the agent can now make her own: because they represent her particularly well, because she is particularly good at them, or because she particularly likes or enjoys them. This is what it means for something to become second nature: the agent makes it part of who she is; and this process itself need not be a reflective one.

So the agent develops a particular, special, relationship with some actions rather than others. The idea of *familiarity* comes in handy here: the agent extends her self and personhood to some of her performances but not to others. And what marks those performances as part of the agent's extended self is not that she thinks about them, that she ponders over them, but the very opposite—that she need not think about them.

“It is a profoundly erroneous truism, repeated by all copy-books and by eminent people making speeches, that we should cultivate the habit of thinking of what we are doing. The precise opposite is the case. Civilization advances by extending the number of operations which we can perform without thinking about them. Operations of thought are like cavalry charges in a battle—they are strictly limited in number, they require fresh horses, and must only be made at decisive moments.” (Whitehead 1911, quoted in Bargh and Chartrand 1999, 464)

I now turn to presenting the causal view of action, focusing primarily on Davidson's (1963) original statement of it: I argue that causalism fails to account for automatic actions.

3. Causalism

Davidson's account of intentional action is put forward in his famous article *Actions, Reasons, and Causes* (1963), where Davidson defends the thesis that reasons explanation (rationalization) is “a species of causal explanation” (3).⁵

⁴ Pippin 2008, 5, attributes a similar point of view to Hegel.

⁵ For anybody who wonders whether Davidson was only interested in his thesis that rational explanations are causal explanations and not also in offering an account of intentional action, the

On Davidson's account, then, some action A is intentional under description ϕ only if that action was caused by a primary reason of the agent comprising of a pro attitude towards actions with a certain property, and a belief that action A, under description ϕ , has that property:⁶

"R is a primary reason why an agent performed the action A, under description d, only if R consists of a pro-attitude of the agent towards actions with a certain property, and a belief of the agent that A, under the description d, has that property." (Davidson 1963, 5)

Pro attitudes, says Davidson, can be "desires, wantings, urges, promptings, and a great variety of moral views, aesthetic principles, economic prejudices, social conventions, and public and private goals and values" (3).

So, on Davidson's account, my flipping the switch is intentional under the description 'flipping the switch' only if it was caused by a primary reason composed of a pro attitude of mine towards actions with a certain property, say the property of 'illuminating the room'; and a belief that my action, under the description 'flipping the switch', has the relevant property of 'illuminating the room'.

The crucial element of Davidson's view is that the primary reason, composed of a pro attitude plus a belief, is the action's cause. As Davidson himself points out (12), causes must be events, but pro attitudes and beliefs are states, and so they cannot be causes. Davidson therefore proposes the "onslaught" (or *onset*, see Lowe 1999, 225) of the relevant mental state as the cause of action.

The difference between a mental state and its onset, which is a mental event, is the same as the difference between believing that there is a bottle on my desk (mental state), and forming the belief (noticing, realizing) that there is a bottle on my desk (mental event). While both kinds of mental states, pro attitude and belief, are always needed—on Davidson's view—to rationalize an action under some description, only one mental event is necessary to cause the action.

As Stoutland emphasizes, the mental states required by Davidson's view must have a very specific content:

"The thesis is a very strong one: it is not saying merely that reasons are causes of behaviour but that an item of behaviour performed for a reason is not intentional under a description unless it is caused by just those reasons whose descriptions yield the description under which the behaviour is intentional. This requires that every item of intentional behaviour have just the right cause." (Stoutland 1985, 46)

following quotes are given as reassurance: "such a reason [...] implies that the action was intentional" (Davidson 1963, 4); "this last point [...] is of interest because it defends the possibility of defining an intentional action as one done for a reason" (4); "to know a primary reason why someone acted as he did is to know an intention with which the action was done." (5)

⁶ Davidson only offers necessary conditions. Any attempt at giving sufficient conditions would, by Davidson's own admission (1973), run against the problem of deviant causal chains—see my last section for a discussion of the issue of deviant causal chains.

So there must be a content relation between the primary reason and the action description in question. Recall Davidson's definition of "primary reason" (Davidson 1963, 5): the belief must make explicit reference to the action description which it rationalizes.

According to Davidson, for example, the following primary reason would not do as a primary reason for flipping the switch: a pro attitude towards 'illuminating the room', and a belief that my action, under description 'turning on the light', has the property of 'illuminating the room'. This primary reason makes no mention of the description 'flipping the switch', and therefore it cannot rationalize my action under the description 'flipping the switch'; even though it will rationalize my action under the description 'turning on the light'.

One note of clarification: the content constraint emphasized by Stoutland is on the belief rather than on the pro attitude. That is to say that, as long as the belief has the 'right' content, the pro attitude can have any content. For example, my action of flipping the switch can be rationalized under the description 'flipping the switch' by a very wide selection of pro attitudes—'turning on the light', 'illuminating the room', 'wasting energy', 'finding some comfort', 'stretching my arm', etc.—as long as the agent believes that her action, under the description in question—'flipping the switch'—has the relevant property towards which the agent has a pro attitude: 'turning on the light', say.

It must be emphasised that causalism does not depend upon endorsing Davidson's Humean reductionism about motivation: many theorists have proposed versions of causalism that appeal, rather, to a single state of *intention* or *plan*.⁷ On these versions of causalism, views will have the following general form: *Sφ* intended intentionally only if *S* intended to *φ*.⁸

4. Causalism and Automatic Actions

Causalism, then, will offer the same account for both automatic and non-automatic actions. It will explain the intentionality of both Sarah's turning a door handle and Sarah's asking for her boyfriend's hand by appealing to those amongst Sarah's psychological states that caused these behaviours. But how does causalism allow for the phenomenology of the two cases being so different? On the one hand, 'turning a door handle' seemingly never really featured in Sarah's thoughts. She never gave it a moment's thought, as they say. And it's no surprise that she wouldn't have: she had much more important things to think about. The thought of marrying Jim, on the other hand, has occupied Sarah's

⁷ See, amongst others, Searle 1983; Brand 1984; Bratman 1984; 1987; Thalberg 1984; Adams and Mele 1989; Bishop 1989; Mele 1992; Mele and Moser 1994.

⁸ This is actually a statement of the so-called Simple View, which not many people endorse (on the Simple View, see Bratman 1984; 1987; Di Nucci 2009; 2010a; McCann 2010; 2011). Other views can be more complicated (see, for example, Bratman's 1987, 119–123, and Mele and Moser's 1994, 253); but this general statement will do for now. In this paper I focus on Davidson, but I have applied those and similar arguments to other causalist approaches elsewhere: Di Nucci 2008; 2009; 2010a; 2011b.

mind for years. Being very methodical, Sarah has pondered the pros and cons many times; she has gone over the same doubts again and again. There has been many a sleepless night when Sarah couldn't think of anything else. And even when Sarah finally resolved to ask, it took her weeks to find the right moment, in the right place. Still, asking Jim to marry her was just one of many intentional actions that Sarah performed during her lifetime; another one of which being her turning the door handle when she went into the living room to finally declare herself to Jim.

Causalists usually dismiss this problem pretty swiftly. Here's Davidson:

“We cannot suppose that whenever an agent acts intentionally he goes through a process of deliberation or reasoning, marshals evidence and principles, and draws conclusions. Nevertheless, if someone acts with an intention, he must have attitudes and beliefs from which, had he been aware of them and had the time, he could have reasoned that his action was desirable (or had some other positive attribute).” (Davidson 1978, 85)

The suggestion is that in automatic cases such as Sarah's turning of the living-room's door-handle, agents aren't normally aware of the psychological states that motivate, cause, and therefore, according to Davidson, rationalize their behaviour. So the idea is that in the non-automatic case, as the result of the painful process of deliberation that Sarah went through, she is very much aware of her reasons. In the door-handle case, on the other hand, her reasons are *phenomenologically silent*. But that does not mean that those reasons, in the shape of psychological states, are not at work: unbeknownst to the agent, they are at their usual *causal* desk, motivating her. So causalism can offer the same story for both automatic and non-automatic actions as long as it can appeal to psychological states being unaware in the automatic case, as Davidson does above. Causalism would then have the virtue of offering a *unified* account of all intentional action.

Mele and Moser offer a similar suggestion:

“A state of having a plan can, however, be dispositional in the way belief is. It can exist even while unmanifested in an event of planning, and even while absent from awareness or consciousness. Laura need not be constantly aware of her plan or continually engaging in *acts* of planning to flip the switch.” (Mele and Moser 1994, 227).

The idea is the same as Davidson's: causalists recognise that it would be implausible to suggest that agents are all the time formulating intentions or plans to perform such routine activities as turning door-handles or flipping switches; or that the psychological states which cause these activities were at the forefront of the agent's mind at the time of acting. Still, causalists suggest, there is an easy

solution at hand: supposing that in all these ‘automatic’ cases, psychological states are dispositional and agents aren’t aware of them.⁹

The causalist strategy is clear: they deploy *phenomenologically silent* dispositional and unaware psychological states in automatic cases. It must be said that not all proponents of such theories have put forward this quick solution to the problem of automatic actions. There are at least two prominent examples of views that appear to take the problem more seriously: Searle (1983, 84–85) distinguishes between *prior intentions* and *intentions in action*, admitting that it would be implausible to suppose that agents form intentions prior to performing every intentional action. On Searle’s account, then, subsidiary and spontaneous actions are not preceded by a *prior intention*. Still, every intentional action, according to Searle, requires an intention in action.¹⁰

In criticizing the so-called *Simple View*, Bratman (1984; 1987) also allows for the fact that not all intentional actions are caused by an intention to do that thing. Still, on Bratman’s *Single Phenomenon View*, there will be an intention in the causal history of all intentional actions. And, furthermore, an action ϕ can be intentional only if it is included in the ‘motivational potential’ of some intention, even if the intention does not necessarily need to be an intention to ϕ (Bratman 1987, 119–120).

I won’t need to go into the details of these alternative proposals because they all have a central feature in common with the general statement of causalism: an action can only be intentional if it is somehow represented either occurrently, dispositionally, or unawarely, in the agent’s mind. If the relevant specific content cannot be traced back to the agent’s mind, then—on all of these theories—the action in question cannot be intentional.

The problem is that, by stipulating that in all automatic cases the relevant psychological states will be unaware (they will *have to* be unaware, as we have seen, to allow for the phenomenological difference between automatic and non-automatic actions), causalism assumes, rather than demonstrating, the correctness of its picture. What reason does the causalist offer us for thinking that, *contra* phenomenological evidence, the relevant psychological states can always be attributed to the agent as the causes of her behaviour?

The causalist merely points out that, since the relevant psychological states can be absent from the agent’s awareness, that the agent does not *appear* to intend to ϕ , and that the agent does not experience herself intending to ϕ cannot

⁹ A note of caution on Mele and Moser’s formulation: it is dangerous to equate, as they do, awareness with consciousness. Since most people, in the philosophy of mind and cognitive science, will accept that all that is required for consciousness is access, then to say that the psychological states that motivate agents in cases of automatic action are ‘absent from consciousness’ actually means denying that the agent has access to them. But if the agent was motivated and caused to act by psychological states to which she has no access, it’s not at all obvious that she has acted intentionally. Psychological states to which agents don’t have access aren’t really up to the agent (not directly, anyway); so that they are rather like psychoanalytic drives or external forces.

¹⁰ We cannot consider Searle’s view a causal one, because Searle does not take the relationship between an intention in action and its action to be a causal relationship, but rather a constitutive one.

be evidence for the claim that in fact the agent does not intend to ϕ . Because there is another story that is compatible with the data: that the agent does not appear to intend to ϕ and does not experience herself intending to ϕ not because she does not indeed intend to ϕ but rather because she is not aware of her intending to ϕ .

But, the reader will have noticed, this is just a draw: the causalist has merely shown that the absence of the relevant mental content from the agent's phenomenology does not imply that the relevant mental content is altogether absent: it could just as well be that the relevant mental content is simply something that the agent is not aware of. But this does not amount to claiming that, indeed, the relevant mental content *is* unaware in automatic cases, because it does not amount to evidence for choosing the conclusion that it is unaware over the conclusion that it is altogether absent.

I am not the only one to have voiced this methodological concern: "Davidson's putative states of believing [...] look like the posits of somebody in the grip of a theory, rather than an independent datum being innocently incorporated into a theory whose correctness is still up for grabs." (Pollard 2006, 9) And here's Dennett:

"Although we are occasionally conscious of performing elaborate practical reasoning, leading to a conclusion about what, all things considered, we ought to do, followed by a conscious decision to do that very thing, and culminating finally in actually doing it, these are relatively rare experiences. Most of our intentional actions are performed without any such preamble, and a good thing, too, because there wouldn't be time. The standard trap is to suppose that the relatively rare cases of conscious practical reasoning are a good model for the rest, the cases in which our intentional actions emerge from processes into which we have no access." (Dennett 1991, 252)

But the causalist *has to* think that the relevant mental content must be there; and that therefore, if it is not overtly there, it must be there somewhere amongst the things the agent is not aware of. Take the case of turning a door-handle again: if all we say is that the agent's plan was to ask her boyfriend to marry her (or, which is the same for present purposes, that the agent's reason for turning the door-handle was that she wanted to ask her boyfriend to marry her), that alone gives us no grounds for thinking that turning a door-handle wasn't an accident or happening, but rather an action—and intentional at that—of an agent.

This point becomes apparent with reference to Davidsonian rationalization: wanting to illuminate the room does not, alone, rationalize flipping the switch. It needs the belief that 'flipping the switch' has the relevant property of 'illuminating the room'. If the agent wanted to illuminate the room and believed that flipping the switch would *not* illuminate the room, then we could not make sense of her behaviour of flipping the switch. So the relevant content-specific belief—the one which actually makes explicit mention of the agent's deed—*must* be there.

And if its presence is not *apparent*, then the agent must be simply *unaware* of the relevant belief.

So the relevant content-specific belief is not merely posited by causalism: it follows from an assumption about the rational coherence of the agent. But now the causalist argument looks suspiciously circular: what are our grounds for thinking that ϕ -ing is intentional? That the agent intended to ϕ . And what are our grounds for thinking that the agent intended to ϕ ? That ϕ -ing is intentional. The same can be said of rationalization: what are the grounds for thinking that it was rationally coherent for E to flip the switch (or: that E's flipping of the switch is rationalized or rationalizable)? That she wanted to illuminate the room and she believed that flipping the switch would illuminate the room. And what are the grounds for thinking that the agent believed that flipping the switch would illuminate the room? That it was rationally coherent for her to flip the switch (or: that E's flipping of the switch is rationalized or rationalizable).

This circularity can be avoided by appealing to two different conceptions of intentionality, rational coherence, or rationalization: an intuitive conception of intentional action that we are trying to analyze into a philosophical definition of intentional action. But the problem is that the causalist is assuming that we cannot do justice to the intuitive conception of intentional action or rational coherence without attributing the relevant psychological states. While the causalist is obviously free to stipulate that her philosophical concept of intentional action requires the relevant psychological states, she cannot also stipulate that the intuitive conception of intentional action requires the relevant psychological states: she would be stipulating that the intuitive conception of intentional action can only be successfully analyzed by her philosophical definition of intentional action—but that is the bit the causalist must argue for. So in the next section I am going to argue, *contra* causalism, that the relevant content-specific psychological states aren't always necessary.

5. Psychological States Aren't Necessary

It isn't at all clear that the attribution of the relevant content-specific belief is actually always necessary. There are different kinds of cases in which the belief isn't necessary: take, for example, my key-ring. There are at least a dozen keys on it, and I can hardly tell them apart. I often find myself in front of my office door, trying out multiple keys to open it. When I try the first key, I intentionally do so. Why am I inserting the key in the key-hole? Because I want to open the door. Do I believe that the key will actually open the door? I don't. But, crucially, I don't believe that inserting the key in the key-hole will interfere with my goal of opening the door either. Had I believed that inserting the key in the key-hole would not open the door, then it would be hard to make sense of my behaviour: why did you insert the key in the key-hole if you thought it was not going to open the door?

But, it will be objected, I must at least believe that it is possible that the key might open the door; otherwise why would I bother? But is that really necessary? It might be that, with experience, I have discovered that the most effective way to open the door is to try each key; that I could indeed tell those keys that might open the door from those keys that won't open the door; but that in order to distinguish them and only try those keys that might open the door it would take too long every time; it is actually quicker to simply try them all. So when I try some of these keys not only don't I have the belief that they will open the door; I might not even have the belief that they *might* open the door. All I believe is that one of these keys will open the door; but that belief doesn't imply a belief that each key might open the door.

So it seems that, in these kinds of cases, all that is required is that the agent does not believe that her action will not satisfy her psychological states; but it is not actually necessary that the agent believes that her action will satisfy her pro attitudes.¹¹ As long as the agent does not actually believe that the action is incompatible with her motives, then she cannot be deemed irrational.

The point is quite general: denying that the agent believes that p or intends to ϕ isn't logically equivalent to claiming that the agent believes $\neg p$ or intends to $\neg \phi$. And while it is apparent that a belief that $\neg p$ or an intention to $\neg \phi$ could not be reconciled with the agent's behaviour without sacrificing the agent's rational coherence, it is not similarly obvious that a belief that p or an intention to ϕ are required: Gert, for example, argues that for an action to be rational, it suffices that it is not irrational. "Defining a rational action simply as an action that is not irrational does not impose a fictitious and misleading uniformity on all rational actions." (Gert 1998, 61) Suppose I am walking to work: I will take most of my steps automatically. Now what I have been arguing is that, in order for an individual step to be an intentional action of mine, I don't need to have a belief that *that* particular step has the relevant property.

It might be objected that I am (mis)attributing to the causalist a view so clearly implausible that it can't really be true to the causalist position. No sensible causalist would suppose that for an individual step in a walk to be an intentional action of mine, I need to have a belief about *that* particular step. So I would be criticizing a figure of straw. But this supposed figure of straw is the position that the causalist is committed to: on the causalist picture, for some step S to be an intentional action of the agent, she needs to have had a pro attitude towards actions with a certain property—say 'going to work'—and a belief that the action in question S satisfies the property of 'going to work'.¹² In the absence of this specific belief—which makes specific reference to the particular action S —the action S cannot be rationalized by the pro attitude alone: why would the agent have performed S if she didn't believe that S would satisfy her goal of going to work? Others might have a different view of rationality (like the aforementioned Gert 1998), but the above is the view the causalist ends up with.

¹¹ For an extended discussion of this topic, see Grice 1971; Bratman 1987 and McCann 1991.

¹² This is also true of whichever version of causalism that requires that there be some reference to the action in question within the agent's psychology.

Here I ought to specify that my argument works differently against reductive causal views than it does against non-reductive causal views.¹³ Still, it is successful against both. The problem for reductive causal views is that, in the absence of the required belief in cases of automatic actions, reductive causal views must allow for some intentional and rationally coherent actions that they cannot rationalize and that therefore, on these views, don't count as intentional—so that automatic actions constitute counterexamples to the necessary condition on intentionality put forward by reductive causal views.

Against non-reductive causal views the argument works differently, because by taking away the intention to ϕ when the agent intentionally ϕ -s, we are taking away the supposed proximal cause of ϕ -ing—so that these views fail because their causal constraint does not hold. This is not the case for reductive causal views because, even without the belief, the cause could be the onset of the relevant pro attitude. So that if I flip the switch because of my pro attitude towards illuminating the room, even without attributing the belief that flipping the switch will illuminate the room, the action's cause could still be the onset of the pro attitude towards illuminating the room.

It could be suggested that there will be available causes in non-reductive accounts too. If the view is, for example, that in order for ϕ to be intentional, I don't actually need an intention to ϕ as long as there is some intention ψ in whose content ϕ -ing is specified, then it could be proposed that the cause of ϕ -ing be the agent's intention to ψ (this is Bratman's Single Phenomenon View, Bratman 1987). But here there is the issue of *contiguity*. If we take, with Hume, "the relation of CONTIGUITY as essential to that of causation" (Hume's *Treatise*, Book I, Part III, Section 2), then it's going to be difficult to always find contiguous intentions, because the whole point of loosening the constraint on intention in the way Bratman proposes is to allow for very general intentions, plans, and policy to explain many subsidiary actions. So that such states of intention will hardly ever be contiguous to the action they are supposed to account for—especially because it is the state's onset that is being proposed as the cause. The same point about the lack of contiguity can also be made for pro attitudes in the case of reductive causalism.

There is another kind of case in which the attribution of the belief does not appear to be necessary: passive action (Frankfurt 1978; Mele 1997; Zhu 2004). Take Frankfurt's famous coasting scenario:

"A driver whose automobile is coasting downhill in virtue of gravitational forces alone might be satisfied with its speed and direction, and so he might never intervene to adjust its movement in any way. This would not show that the movement of the automobile did not occur under his guidance. What counts is that he was prepared to intervene if necessary, and that he was in a position to do so more or less effectively. Similarly, the causal mechanisms which stand ready

¹³ To be sure: on reductive causal views such as Davidson's intentions are reducible to a belief-desire pair; non-reductive causal views such as Bratman's deny this.

to affect the course of a bodily movement may never have occasion to do so; for no negative feedback of the sort that would trigger their compensatory activity might occur. The behaviour is purposive not because it results from causes of a certain kind, but because it would be affected by certain causes if the accomplishment of its course were to be jeopardized.” (Frankfurt 1978, 160)

The driver is intentionally, but passively, coasting downhill. Does she need an intention to coast or a belief that coasting will satisfy her motives? Not really: we could as easily suppose that the reason why the driver does not intervene is that she is not dissatisfied with her coasting. We don't need to suppose that the driver actually wants or intends to coast in order to make rational sense of her behaviour. Here, again, were the driver to believe that coasting was incompatible with her goals, then her behaviour would have to be deemed irrational. But as long as the driver does not believe that, we can make perfect sense of it. We could certainly suppose that the driver intends to coast; but the crucial point is that we don't need to suppose so (Zhu 2004 also makes this point).¹⁴

Finally, there is a third case where the attribution of the relevant belief or intention does not appear possible: those are cases of so-called 'arational actions', first presented by Hursthouse (1991). Here are some of her examples: “rumpling someone's hair, “throwing an 'uncooperative' tin opener on the ground” (58), jumping up and down in excitement, “covering one's face in the dark [out of shame]” (ibid.), “covering one's eyes [in horror] when they are already shut” (ibid.). These actions appear to be intuitively intentional, but we can't really rationalize them through a pair of the agent's psychological states. This case, though, might be only relevant for Davidsonian rationalization and not for causal views whose posits are simply intentions. If we drop the rationalization requirement, there is nothing against attributing the intention to, say, rumple someone's hair—apart from, possibly, the sudden and emotional character of many examples of arational actions.

So content-specific psychological states aren't always necessary. Therefore we need not conclude that, in cases where there appear to be no such preceding psychological states, these psychological states must then be unaware. Here there are at least two ways out for causalism: one strategy for a defender of causalism would be to argue that, in all cases of intentional action there are independent grounds for the attribution of the relevant content-specific psychological states. Grounds that are independent, that is, from the idea that I have been criticising: that without attributing the relevant psychological states, we couldn't make sense of the agent's behaviour. What would count as independent grounds, for example, is the agent's phenomenology and experience of acting—but, as we have seen, that cannot be appealed to in the case of automatic actions.

Also, the grounds for the attribution of these psychological states need to be independent from the attribution to the agent of the relevant action: causalists

¹⁴ See also Sartorio on omissions (2005; 2009; on this, see also Clarke 2010; 2011). I have discussed Frankfurt at length elsewhere: Di Nucci 2010b; 2011a and 2011b.

don't want the attribution of these psychological states to be, in the words of McDowell (1978, 79), *consequential* to their taking the agent to act in a certain way. Because then the psychological state wouldn't be "independently intelligible" (84), but rather conceptually dependent upon the action it is supposed to explain. And then the causalists' old foes (Anscombe 1957; Hampshire 1959; Melden 1961) would be vindicated: the relationship between reason (or intention) and action would indeed be conceptual, so that the requirement that the cause be independent from the effect would be violated. In the next section I am going to dismiss some other potential independent grounds for the attribution of the relevant content-specific psychological states.

An alternative strategy for a supporter of causalism could be here to argue that a sufficient reason for endorsing causalism (and therefore a sufficient reason for thinking that the relevant content-specific psychological states are always unaware when they are not apparent) is that there is no coherent alternative available. In the section after next I am going to suggest an alternative to causalism.

6. Self-Attributions

One independent argument for the attribution of the relevant psychological states could be that agents often self-attribute (and, crucially, would always self-attribute if challenged) an explanation that either features the relevant psychological states or depends upon those psychological states.

The idea, then, is that the fact that agents themselves would self-attribute primary reasons (or parts thereof) or intentions even in automatic cases provides us with an argument for always attributing the required unaware psychological states. So, even though agents might have been unaware of their beliefs or intentions, the fact that they can report them afterwards is a reason to think that the agent had the required content-specific psychological state.

But can we accept the agent's own version of events? If we take the agent's own explanation of her behaviour as good evidence, then we have to accept a sort of epistemic authority of the agent over her reasons and actions. And the philosopher, at least as much as the layman, cannot just accept what people say about themselves (Tanney 1995, 10, puts this point rather nicely). The concern here is a methodological one: a theory that need not rely on the agent being both truthful and correct about herself is a methodologically superior theory (Pollard 2005 makes a similar point).

It is not only both philosophical and lay common sense that speak against relying on self-attributions. It is also psychoanalysis, which tells us that agents are often mistaken (in denial) about their reasons, in the way of both being ignorant of one's actual reasons, and mistaking other considerations for one's actual reasons. Both these kinds of self-deceptions tell us not to trust self-attributions.

Another reason for being sceptical about agents' self-attributions is that, when they come in the shape of answers to questions, they might be influenced

by the way in which the question has been asked: so that a Davidsonian question might *lead* a Davidsonian answer. The agent might provide a primary reason only because the question assumed one; but they might have not actually volunteered one. When asked ‘Did you intend to ϕ ?’ most agents will, in normal circumstances, answer positively; but all they mean might just be ‘I certainly didn’t intend to not ϕ ’: and the difference here is important. The latter does not point to the presence of an actual mental state. But the causalist needs actual mental states because she needs causes.

One more reason not to accept self-attributions as evidence comes from everyday language. People are often accused of, rather than congratulated for, ‘rationalizing’ their behaviour. The accusation is that they make up their reasons *ad hoc* to make sense of their behaviour. Suppose the meal you’ve just cooked for your friends turns out tasteless because you have forgotten salt altogether. When the complete absence of salt is pointed out to you, you might rationalize your behaviour to your friends by citing health concerns. What has actually happened is that you forgot. What you are doing, there, is constructing a story that will make sense of your actions and get you off the hook: you are trying to avoid responsibility and look good. Concern for your friends’ health would rationalize, in Davidson’s sense, your actions if you had actually acted from those considerations. But since you didn’t, those considerations do not rationalize your behaviour. What you are actually doing, in making up your story—‘rationalizing’ in the everyday language sense—is, in short, lying.¹⁵

Note, also, the similarity between the way in which everyday language understands the practice of rationalizing and my general line of argument against causalism: both in everyday language and in my argument the respective practices that go under the name of ‘rationalization’ are accused of constructing, rather than reporting, reasons. And a braver philosopher than I am would claim that to be in itself an argument against causalism.

The crucial point, indeed, is that constructions are not good enough for causalism. Causalism needs descriptions because it needs actual mental states and actual causes. It has been argued that constructions could be good enough to make sense of an agent’s behaviour as rational (see Pollard 2003, 424). But as long as those constructions would not point to the agent’s psychology (her mental states), then they would not be good enough for causalism.

So the agent’s self-attributions of reasons don’t seem to be good enough evidence for thinking that the agent actually acted from the reasons or intentions she attributes to herself. Should we then conclude that self-attributions should be disregarded altogether as a guide for understanding an agent’s reasons or intentions? Such a conclusion seems too strong: all I have shown is that, when our only ground for concluding that an agent acted for some reason is that the agent thinks or says that she acted for that reason, then we do not have sufficient

¹⁵ The point I am making here is one of method: I am obviously not saying that agents are always or even often in denial about their reasons, or that they always or even often ‘rationalize’ in its street-sense. I am only saying that because these phenomena are pretty common, a theory that does not rely on self-attributions is a methodologically superior theory to one which does.

grounds for attributing that reason to the agent. Obviously, in normal cases we will accept the agent's version; but that is just because, in normal cases, we will have other elements which substantiate that version (environment, circumstances, what we know about the agent's past, habits, and preferences, what we know about human nature, other people's versions, etc.).

There is another problem with self-attributions that is independent of whether or not we are warranted in accepting them as evidence: that agents might very well truthfully and correctly self-attribute reasons, but that does not mean that they are pointing to some actual mental states of theirs as causes of their actions.

The point is two-fold: it means, on the one hand, that self-attributions might lend support to the Humean belief-desire model of motivation (according to which both beliefs and desires are necessary to motivate an agent to act, see Smith 1987; 1996), but that does not mean that they lend support to reductive causalism, according to which not only is the belief-desire pair necessary to motivate the agent to act, but actually the belief-desire pair causes the agent's actions. So even if agents did self-attribute reasons, that would not mean that they were self-attributing causes—and this point applies to non-reductive causalism too, even though that view might be constructed independently of a Humean view of motivation. This point is pretty simple: on the assumption that a Humean need not accept Davidson's (1963) causal thesis, then self-attributions do not support causalism. Because they do not explicitly point to causes and, on the above assumption, what they point to—reasons—do not necessarily need to be causes.

Secondly, the fact that agents self-attribute reasons does not mean that they necessarily self-attribute mental states. Indeed, that would be assuming internalism about reasons for action. According to externalists such as Stout (1996), Collins (1997), and Dancy (2000), reasons are not psychological states of the agent, but facts of the mind-independent world in the light of which agents act.

7. Guidance

Here I am going to sketch an alternative to causalism that can account for the intentionality of automatic actions because it does not rely on preceding psychological states as the causes of actions.

The general suggestion is that, instead of founding our account of intentional action on the psychological states which motivate the agent, we rather focus on whether or not the agent had control over her behaviour (as originally suggested by Frankfurt (1978): some bodily movement is an action only if it is under the agent's guidance). This would have multiple advantages over the causal view: it could account for the intentionality of automatic actions, for a start; because the fact that agents don't think about their automatic performances does not suggest that they might have reduced or no control over them. Indeed, the opposite is more likely: an agent's familiarity with her automatic actions and habits means

that she need not consciously and continuously monitor her performance. She is so good at it that she no longer needs to pay attention.

A control-based view would also have the advantage of being able to accommodate those cases presented above, passive actions and arational actions. And, perhaps most importantly, a control-based view wouldn't be challenged by deviant causal chains. This famous kind of counterexample to causalism derives its strength exactly from the fact that, despite these scenarios meeting the causalist conditions for intentional action, what happens is obviously not an intentional action of the agent because the agent has no control over it.¹⁶

The paradigmatic deviant case was set by Davidson in *Freedom to Act* with the climber example: suppose a climber decides to rid herself of the "weight and danger of holding another man on the rope" (Davidson 1973, 79).¹⁷ The decision to commit such a horrible act unnerves the climber so much that she loosens her grip on the rope, thereby ridding herself of the other man on the rope. Her decision to rid herself of the other climber both causes and rationalizes the loosening of the rope; but the agent did not intentionally loosen the rope: it was an accident.

Trying to solve the problem of deviant causal chains has kept causalists busy for decades: Mele and Moser have proposed that psychological states should be constructed as "sustaining and guiding actions" (Mele and Moser 1994, 236); Thalberg that causation should be understood as "on-going, continuous, or sustained" (Thalberg 1984, 257).¹⁸ I can't address here the relative merits of all such proposals (for a counterexample to these proposed amendments that I take to be successful, see Moya 1990); I am only going to emphasize that a control-based view wouldn't be subject to the problem of deviant causal chains, because if control is necessary for intentional action, deviant scenarios won't count as cases of intentional action because they are characterized exactly by the lack of control.

This is not the place to spell out my control-based view in much detail (this paper is *only* a critique of causalism), so I will simply state my view and point to the way in which it is different from Frankfurt's original proposal.¹⁹ On Frankfurt's view, some bodily movement counts as an action only if it is under some agent's guidance (Frankfurt 1978, 157). A movement is under an agent's guidance if the agent is in a position to make compensatory adjustments if necessary.

Still, for a bodily movement to count as an intentional action, the agent, according to Frankfurt, must have intended to perform it.

¹⁶ Just to clarify: deviant causal chains are a problem for causalism only if causalism wants to offer both necessary and sufficient conditions for intentional action; if causalism only offers necessary conditions, then deviant causal chains are no counterexample. But then causalism would be offering a partial and incomplete account of intentional action—while my guidance view promises a full-blown account.

¹⁷ The other famous deviant scenario is the one in which Fred runs over his uncle by accident on his way to kill his uncle (Chisholm 1966).

¹⁸ Similar proposals can be found in Brand 1984; Alston 1986; Audi 1986; Bishop 1989; Mele and Adams 1989; Mele 1992.

¹⁹ More details can be found in Di Nucci 2008 and 2011b.

“The term ‘intentional action’ may be used, or rather mis-used, simply to convey that an action is necessarily a movement whose course is under an agent’s guidance. When it is used in this way, the term is pleonastic. In a more appropriate usage, it refers to actions which are undertaken more or less deliberately or self-consciously—that is, to actions which the agent intends to perform. In this sense, actions are not necessarily intentional.” (Frankfurt 1978, 159)

So for the concept of ‘intentional action’ Frankfurt resorts again to intentions. I think, on the other hand, that we can offer an account in which guidance is sufficient for intentional action, as long as we can find a way to distinguish between intentional and *unintentional* actions solely through the concept of guidance. I think the following ‘guidance’ view does the job: *E* ϕ -s intentionally iff ϕ -ing is under *E*’s guidance, where ϕ -ing is under *E*’s guidance only if the agent can reasonably be expected to know or find out about her movements under the relevant description, ϕ -ing.

So that if Tom’s phone number has been assigned to Tim overnight, and when I phone Tom in the morning, I end up speaking to Tim, I cannot be said to have intentionally phoned Tim, because I couldn’t have reasonably been expected to know or find out that Tom’s phone number had been reassigned to Tim. So I phone Tim unintentionally—even though the movements which constitute my phoning Tim are under my guidance, these movements resulting in the action-description ‘phoning Tim’ is out of my *rational* control, because I cannot reasonably be expected to know that my movements will result in my phoning Tim, nor can I reasonably be expected to check, every morning, that all of my friends still have the phone numbers they had the night before.

Without going into too much detail, here I am going to touch upon two important aspects of this view: its intuitiveness and the reasonable expectations component. With the latter, the idea is that what the agent can reasonably be expected to know or find out is context-dependent: how important my phone call to Tom is will make a difference to what it is reasonable to expect of me in terms of making sure that I connect to Tom and not to someone else when I try ringing him. If I just want to chat is one thing; if I or someone else needs urgent medical attention is an altogether different one.

Similarly, the standards of reasonableness will also be somewhat relative to the agent in question and what kinds of habits, attitudes, and history the agent has with relation to phoning people: does she always check the number? Does she often appeal to memory alone? How many times has she dialled that number before? These sorts of considerations will partially determine what can reasonably be expected of the agent.

Here I should emphasize that relativizing to contexts and agents does not mean relativizing back to psychological states, which would make my guidance view depend on the agent’s mind just like causal views: in the scenario above, for example, there is no mention of the agent’s beliefs and desires.

There is also the question of whether my guidance view fits with our common conception of intentional action. Suppose that Tom, the evening before, had informed me that his number was going to be reassigned to Tim. Suppose that I took note of it, editing my phonebook accordingly. Still, in the morning I dial Tom's old number out of habit and I'm connected to Tim. It was reasonable to expect that I knew that I was actually phoning Tim. Still, can we say that I phoned Tim intentionally? I didn't mean to.

I think that there are times when epistemic and rational pressures are such that it makes little sense to speak of not having done something intentionally. All it does is describing, again, the agent's motivation. But what is at issue is, rather, her action. And sometimes what the agent doesn't know or believe is so striking that it overshadows what she does know and believe. Sometimes what we ought to have known rings so loud that what we did know hardly matters. And it is agents themselves who recognise this: that's why Oedipus blinds himself; whether or not he *wanted* to kill his father and marry his mother doesn't matter to him then. Oedipus isn't cursing his reasons, but his deeds. We ought to be mindful of that when defining intentional action.

8. Conclusion

I have shown that so-called automatic actions—performances that we complete routinely without thinking, such as turning a door-handle—pose a challenge to causalism because they do not appear to be preceded by the relevant psychological states. Causalism cannot prove, I have argued, that the agent is simply unaware of the relevant psychological states in all cases of automatic actions. This is because, in the first place, we can often make coherent sense of the agent's behaviour without necessarily appealing to the relevant content-specific psychological states. Secondly, there are no obvious independent arguments in favour of the attribution of the relevant psychological states in automatic cases: agents' own self-attributions, I have argued, won't do as evidence. Finally I have put forward an alternative view of intentional action, inspired by Frankfurt, based on the concept of 'guidance'.

References

- Adams, F. and A. Mele (1989), "The Role of Intention in Intentional Action", *Canadian Journal of Philosophy* 19, 511–31.
- Alston, W. (1986), "An Action-Plan Interpretation of Purposive Explanations of Actions", *Theory and Decision* 20, 275–299.
- Anscombe, G. E. M. (1957), *Intention*, Oxford: Basil Blackwell.
- Aristotle (1925), *Nicomachean Ethics*, translated with an introduction by David Ross, Oxford: Oxford University Press.

- Audi, R. (1986), "Acting for Reason", *Philosophical Review* 95, 511–546.
- Bargh, J. A. (2006), "What Have We Been Priming All These Years?", *European Journal of Social Psychology* 36, 147–168.
- (2008), "Free Will Is Un-natural", in: Baer, J., J. C. Kaufman and R. F. Baumeister (eds.), *Are We Free?*, Oxford: Oxford University Press, 128–154.
- , M. Chen and L. Burrows (1996), "Automaticity of Social Behavior: Direct Effects of Trait Construct and Stereotype Activation on Action", *Journal of Personality and Social Psychology* 71, 230–244.
- and T. L. Chartrand (1999), "The Unbearable Automaticity of Being", *American Psychologist* 54, 462–479.
- Bishop, J. (1989), *Natural Agency. An Essay on the Causal Theory of Action*, Cambridge: Cambridge University Press.
- Brand, M. (1984), *Intending and Acting*, Cambridge/MA: The MIT Press.
- Bratman, M. (1984), "Two Faces of Intention", *Philosophical Review* 93, 375–405.
- (1987), *Intention, Plans, and Practical Reason*, Cambridge/MA: Harvard University Press.
- Chisholm, R. (1966), "Freedom and Action", in: Lehrer, K. (ed.), *Freedom and Determinism*, New York: Random House, 11–44.
- Clarke, R. (2010), "Intentional Omissions", *Nous* 44(1), 158–177.
- (2011), "Omissions, Responsibility, and Symmetry", *Philosophy and Phenomenological Research* 82(3), 594–624.
- Collins, A. W. (1997), "The Psychological Reality of Reasons", *Ratio* X, 108–123.
- Dancy, J. (2000), *Practical Reality*, Oxford: Oxford University Press.
- Davidson, D. (1963), "Actions, Reasons, and Causes", *Journal of Philosophy* 60, 685–700.
- (1971), "Agency", in: Binkley, R., R. Bronaugh and A. Marras (eds.), *Agent, Action, and Reason*, Toronto: University of Toronto Press, 3–37.
- (1973), "Freedom to Act", in: Honderich, T. (ed.), *Essays on Freedom and Action*, London: Routledge and Kegan Paul, 137–56.
- (1978), "Intending", in: Yovel, Y. (ed.), *Philosophy of History and Action*, Jerusalem: The Magnes Press, The Hebrew University, 83–102.
- (1980), *Essays on Actions and Events*, Oxford: Oxford University Press.
- Dennett, D. C. (1991), *Consciousness Explained*, London: Penguin.
- Di Nucci, E. (2008), *Mind Out of Action*, Saarbrücken: VDM Verlag.
- (2009), "Simply, False", *Analysis* 69(1), 69–78.
- (2010a), "Rational Constraints and the Simple View", *Analysis* 70(3), 481–486.
- (2010b), "Refuting a Frankfurtian Objection to Frankfurt-type Counterexamples", *Ethical Theory and Moral Practice* 13(2), 207–213.
- (2011a), "Frankfurt Counterexample Defended", *Analysis* 71(1), 102–104.
- (2011b), "Frankfurt versus Frankfurt: A New Anti-causalist Dawn", *Philosophical Explorations* 14(1), 117–131.
- (forthcoming), "Priming Effects and Free Will", *International Journal of Philosophical Studies*.

- Dreyfus, H. (1988), "The Socratic and Platonic Bases of Cognitivism", *AI & Society* 2, 99–112.
- (2005), "Overcoming the Myth of the Mental", *Topoi* 25(1–2), 43–49.
- and S. Dreyfus (1984), "Skilled Behavior: The Limits of Intentional Analysis", in: Lester, E. (ed.), *Phenomenological Essays in Memory of Aron Gurwitsch*, Washington: The University Press of America, 23–39.
- Frankfurt, H. (1978), "The Problem of Action", *American Philosophical Quarterly* 15, 157–162.
- Gert, B. (1998), *Morality: Its Nature and Justification*, Oxford: Oxford University Press.
- Ginet, C. (1990), *On Action*, Cambridge: Cambridge University Press.
- Grice, H. P. (1971), "Intention and Uncertainty", *Proceedings of the British Academy* 57, 263–79.
- Hampshire, S. (1959), *Thought and Action*, London: Chatto and Windus.
- Hornsby, J. (1980), *Actions*, London: Routledge & Kegan Paul.
- Hume, D. (1978), *Treatise of Human Nature*, Nidditch edition, Oxford: Oxford University Press.
- Hursthouse, R. (1991), "Arational Actions", *Journal of Philosophy* 88(2), 57–68.
- James, W. (1890), *The Principles of Psychology*, London: Macmillan.
- Lowe, J. (1999), "Self, Agency, and Mental Causation", *Journal of Consciousness Studies* 6, 225–39.
- McCann, H. (1991), "Settled Objectives and Rational Constraints", *American Philosophical Quarterly* 28, 25–36.
- (2010), "Di Nucci on the Simple View", *Analysis* 70(1), 53–59.
- (2011), "The Simple View Again: A Brief Rejoinder", *Analysis* 71(2), 293–295.
- McDowell, J. (1978), "Are Moral Requirements Hypothetical Imperatives?", *Proceedings of the Aristotelian Society*, Supplementary Volume 52, 13–29.
- (1996[1994]), *Mind and World* (with a new introduction by the author), Cambridge/MA: Harvard University Press.
- Melden, I. (1961), *Free Action*, London: Routledge Kegan & Paul.
- Mele, A. (1992), *Springs of Action*, Oxford: Oxford University Press.
- (1997), "Passive Action", in: Holmstrom-Hintikka, G. and R. Tuomela (eds.), *Contemporary Action Theory*, Dordrecht: Kluwer, 135–143.
- and P. K. Moser (1994), "Intentional Action", *Nous* 28, 39–68.
- Moya, C. (1990), *The Philosophy of Action*, Cambridge: Polity Press.
- O'Shaughnessy, B. (1980), *The Will: Volume 2, A Dual Aspect Theory*, Cambridge: Cambridge University Press.
- Pippin, R. B. (2008), *Hegel's Practical Philosophy*, Cambridge: Cambridge University Press.
- Pollard, B. (2003), "Can Virtuous Actions Be Both Habitual and Rational?", *Ethical Theory and Moral Practice* 6, 411–425.
- (2005), "Naturalizing the Space of Reasons", *International Journal of Philosophical Studies* 13(1), 69–82.

- Pollard, B. (2006), "Actions, Habits, and Constitution", *Ratio* 19, 229–248.
- Rietveld, E. (2008), "Situated Normativity: The Normative Aspect of Embodied Cognition in Unreflective Action", *Mind* 117, 973–1001.
- Ryle, G. (1949), *The Concept of Mind*, London: Penguin.
- Sartorio, C. (2005), "A New Asymmetry between Actions and Omissions", *Nous* 39, 460–482.
- (2009), "Omissions and Causalism", *Nous* 43, 513–530.
- Searle, J. (1983), *Intentionality*, Cambridge: Cambridge University Press.
- Smith, M. (1987), "The Humean Theory of Motivation", *Mind* 96, 36–61.
- (1996), *The Moral Problem*, Cambridge/MA: Harvard University Press.
- Stout, R. (1996), *Things That Happen Because They Should*, Oxford: Oxford University Press.
- Stoutland, F. (1985), "Davidson on Intentional Behaviour", in: LePore, E. and B. P. McLaughlin (eds.), *Actions and Events*, Oxford: Basil Blackwell, 44–9.
- Tanney, J. (1995), "Why Reasons May Not Be Causes", *Mind & Language* 10, 103–126.
- Thalberg, I. (1977), *Perception, Emotion, and Action*, Oxford: Basil Blackwell.
- (1984), "Do Our Intentions Cause Our Intentional Actions?", *American Philosophical Quarterly* 21, 249–260.
- Thomson, J. J. (1977), *Acts and Other Events*, Ithaca: Cornell University Press.
- Withehead, A. N. (1911), *An Introduction to Mathematics*, New York: Holt.
- Wittgenstein, L. (1969), *On Certainty*, Oxford: Basil Blackwell.
- Zhu, J. (2004), "Passive Action and Causalism", *Philosophical Studies* 119, 295–314.