

INTENTIONS AND THE REASONS FOR WHICH WE ACT

ULRIKE HEUER

Many of the things we do in the course of a day we don't do intentionally: blushing, sneezing, breathing, blinking, smiling—to name but a few. But we also do act intentionally, and often when we do we act for reasons. Whether we always act for reasons when we act intentionally is controversial. But at least the converse is generally accepted: when we act for reasons we always act intentionally. Necessarily, it seems. In this paper, I argue that acting intentionally is not in all cases acting for a reason. Instead, intentional agency involves a specific kind of control. Having this kind of control makes it possible to modify one's action in the light of reasons. Intentional agency opens the possibility of acting in the light of reasons. I also explain why when we act with an intention (and not just intentionally in a broader sense) we act for reasons. In the second part of the paper, I draw on these results to show that the dominant view of reasons to intend and the rationality of intentions should be rejected.

Many of the things we do in the course of a day we don't do intentionally: blushing, sneezing, breathing, blinking, smiling—to name but a few. But we also do act intentionally, and often when we do we act for reasons. Whether we always act for reasons when we act intentionally is controversial. But at least the converse is generally accepted: when we act for reasons we always act intentionally. Necessarily, it seems (Setiya 2004, 2007). What is it to act intentionally? And how does acting intentionally relate to acting for reasons? These are my guiding questions for §§i–vi. In §§vii–x I will put the results to work, and argue that the perhaps dominant view of reasons to intend and the rationality of intentions should be rejected.

Is Acting Intentionally Acting for a Reason? According to some influential approaches to understanding intentional agency the answer is yes (Davidson 1963; Hornsby 1993). What sets intentional behaviour apart from those other things we do, but don't do intentionally, is that it is acting for reasons.

As Elisabeth Anscombe (1963) sees it, when a person acts intentionally, we can ask her 'why?' (the question has application, as Anscombe puts it), and she will typically answer it by giving her reasons for so acting. The why-question is ambiguous, though. It is perfectly applicable to non-intentional behaviour. When someone slips on a freshly polished floor, we can ask 'Why did you slip?'. 'Because the floor was wet' might be the answer. But slipping isn't intentional. Thus, if Anscombe's why-question has application only if a person acts intentionally, the question must be understood implicitly as aiming at a particular kind of answer. Anscombe's question, it seems, is a question that inquires about the point or the value that the person sees in acting as she does: her reasons for the action. When she acts intentionally she will offer in reply (what she believes to be) a normative reason for her action.² So if, for instance, the question 'Why did you slip?' would be answered by (what the agent takes to be) a normative reason, it would indeed show that the action is intentional. An actor may answer, 'I slipped because it was part of the play's directions'. If the actor slipped for that reason, she slipped intentionally. That an action

¹ The reasons for which a person acts may not be the reasons for which she should have acted. She could be irrational (e.g. weak-willed), but she might also simply be mistaken. In that case there may be no reason to act as she does; she only believes there is. The qualification in parenthesis is meant to make room for this possibility. The agent's reasons should be understood throughout in this way, allowing for mistaken beliefs.

² I will simply assume here that reasons to act are normative reasons, and that the reasons for which someone acts are seen by her as normative reasons (rightly or wrongly). This is controversial: in particular, Kieran Setiya (2007, pt. 1, ch. 6) contests the view. As he sees it, the reasons for which we act needn't be reasons that we regard as normative reasons — or, to put it differently, we don't always act under the guise of the good. Discussing his view would take me too far afield here.

is done for a normative reason³ seems a sufficient condition for it to be intentional.⁴ Thus Anscombe's view easily accommodates the generally accepted truth that actions done for reasons are necessarily intentional.

But while acting for a reason is a core case of acting intentionally, there are other forms of intentional agency that Anscombe's view cannot accommodate. They show that acting for a reason is not a necessary condition of acting intentionally. Anscombe was presumably not concerned with establishing necessary and sufficient conditions for acting intentionally, but the cases that escape her account are sufficiently important to call for an explanation of what makes them intentional. The cases:

Frankfurt's Spider. Harry Frankfurt invites us to 'consider the difference between what goes on when a spider moves its legs when making its way along the ground, and what goes on when its legs move in similar patterns and with similar effect because they are manipulated by a boy who has managed to tie strings to them. In the first case the movements ... are attributable to the spider, who makes them. In the second case the same movements occur but they are not made by the spider, to whom they merely happen. This contrast ... parallels the more familiar contrast between the sort of event that occurs when a person raises his arm and the sort that occurs when his arm goes up without his raising it' (Frankfurt 1978, p. 78). And perhaps we could add: the contrast also parallels that between the actor who slips

³ I take it that normative reasons are properties of an action that make it good. Mistakenly believed to be normative reasons may be properties of an action of which the agent falsely believes that they make it good. Or they may be good-making features that the agent wrongly believes her action has. There is more than one way of being mistaken.

⁴ The problem of 'deviant causal chains' is relevant here: the actor while attempting to take a position that allows her to slip convincingly may lose her balance and slip. The slipping would have been triggered by her reason, but it would not have been intentional. However, I take it that Setiya's reply (2007, pp. 56 ff.) solves this problem: when a person acts for a reason the reason does not just trigger the action, but guides it throughout. Deviant causation doesn't require rejecting the causal theory of action, but only the view that causes are triggers.

because it is required by her part and the person who slips because the floor is wet. But if so, how could the fact that the actor had a reason for slipping explain the difference? Surely spiders don't crawl because they believe that there is a normative reason for them to do so.

Expressive Actions. Rosalind Hursthouse imagines a number of actions done out of anger like 'throwing an "uncooperative" tin opener on the ground or out of the window, kicking doors that refuse to shut and cars that refuse to start, tying towels that keep falling off a slippery towel rail on to it very tightly and then consolidating the knots with water; muttering vindictively "I'll show you" ...' (Hursthouse 1991, p. 58). These actions express emotions, but they are not done in order to do so, or for the reason that they release the anger. They don't seem to be done for reasons, yet they are intentional.

Doodling. Finally, certain things we do when passing the time are intentional, but not done for reasons.⁵ Doodling while listening to a philosophy paper is intentional behaviour, but we don't normally doodle for a reason.

Thus the first stab at an answer to our guiding question fails. The answer to the question 'What is it to act intentionally?' is not that it is acting for a reason. Anscombe herself grants that one admissible answer to the why-question is 'No reason' (1963, §§17–18). Expressive actions and doodling illustrate this possibility. But allowing for it seems puzzling if the view Anscombe set out to defend were that intentional actions are done for reasons. I will therefore understand Anscombe as claiming that acting for a reason is sufficient for acting intentionally, and inquiring about a person's reasons is often the best way of finding out what it is that she does intentionally. However, we have yet to establish what those actions that are not done for a reason, but are intentional, have in common

⁵ Joseph Raz (2011, ch. 4) uses examples of this kind to draw the distinction between acting intentionally and acting with an intention.

with those that are done for reasons. What is it that distinguishes both of them from unintentional behaviour?

II

The Unity of Intentional Agency. Is there a unifying feature which is shared by those intentional actions that Anscombe focuses on and the intentional actions in the three examples? What explains that they all are instances of intentional agency? An answer that may seem obvious is that intentional behaviour involves intentions — that is its unifying feature (Setiya 2007; Bratman 1987).

According to the Simple View, as Michael Bratman calls it, when a person ϕ s intentionally, she intends to ϕ . But our counterexamples to the claim that intentional actions are done for reasons are also counterexamples to this view of intentional agency: spiders don't have intentions; expressions of emotions are not done with an intention to express the emotion (except perhaps when the agent has been prompted by her therapist not to bottle up her emotions); and we don't doodle with the intention of doing so.

Bratman himself rejects the Simple View and offers an alternative.⁶ As he sees it, the Simple View is a special case of the view that is actually correct, the Single Phenomenon View: 'On this more general view, intentional action and the state of intention both involve a certain common state, and it is the relation of an action to this state that makes that action intentional' (Bratman 1987, p. 112). The common state is intention: according to Bratman, a person intentionally ϕ s only when she acts with some intention, but it needn't be the intention to ϕ . For example, when I intentionally set out to do a hundred push-ups (something I know I cannot do), I needn't intend to do a hundred push-ups, but I at least have to intend to try to do so. So what I do is intentional because it involves the intention to try.

⁶ I will not discuss here Bratman's famous video game example, which he uses to make this point (1987, ch. 8).

However, at least two of our three examples are counterexamples to the Single Phenomenon View as well. Some expressive actions may be amenable to Bratman's suggestion: while they don't involve intentions to express the emotion, they may involve intentions to (say) throw the tin opener on the floor, or to tie the towel to the rail. But there are also expressive actions which Bratman's suggestion cannot accommodate: banging the table in frustration, say. Furthermore, the suggestion does not work for doodling (and other ways of passing the time). There is no intention to doodle, and presumably no other intention either—certainly no intention to try. I will from now on dispense with the spider, simply because I know too little about the ways in which spiders move. However, it is important to bear in mind that at least some animal actions are intentional, whether or not this is true of spiders.

What then makes intentional behaviour intentional? Not that it is done for a reason. Not that it is done with an intention. Not even that the agent must have some intention in doing what she does.

III

Intentional Control. Another common, but perhaps less explored answer is that intentional actions must be under the agent's control. After all, what we do unintentionally (e.g. slip on the wet floor) is often not under our control. But there is an immediate objection. Some of the things we don't do intentionally are also under our control: we can control our breathing or smiling to some degree. We can suppress a smile, or modify our breathing. Others are not: while both breathing and the movements of the digestive system are not intentional, only the former is (to some degree) under our direct control. So the ability to control draws the line in the wrong place. But even if our ordinary smiling and breathing is not intentional, suppressing a smile or taking a deep breath may well be intentional. So perhaps that we *can* control behaviour does not make it intentional, but when we *actually* control it, it is intentional.

However, there also seems to be *actual control* which is not intentional: I control the various movements of my limbs when I walk or cycle, but I do not control them intentionally. I walk and cycle intentionally, but I do not contract my muscles intentionally. Yet it seems that I do control these movements: I modify them when the circumstances change as when I cycle uphill, or walk on rough or slippery ground. Perhaps highly skilled performances are even more striking in this regard: a violinist or a surgeon controls the very precise movements of her fingers, but they are intentional only at a general level: the violinist intentionally plays the second movement of Bach's Violin Concerto no. 2, say. While playing, she controls each and every motion that she makes in the course of the performance, but not every movement of her fingers is therefore (qua finger movement) intentional.⁷ Is control as a mark of intentional agency another blind alley then?

We may be able to distinguish between two kinds of control: the violinist's control over the movements of her fingers, and a more explicit kind of control she has over exercising her skill in the first place. The latter involves knowledge or belief of a certain kind: when playing the concerto, the violinist believes that this is what she is doing. More precisely, she has a self-referential belief that she is playing the concerto. The control of the movements of the fingers may also involve knowledge, since it is the result of learning, experience and training. A newborn baby, for instance, cannot control her limb movements at all. She has to learn how to do this. The knowledge that the baby acquires in learning to control her movements is of a different kind from the propositional knowledge that the violinist has when she intentionally performs the concerto: it is non-propositional knowing how.⁸ I don't mean to say that knowing how is always

⁷ The precise movements may be guided in some way by the intention: they are adjusted and modified so as to achieve their purpose. For an illuminating discussion of the various kinds of control that are involved in skilled behaviour, and its relation to the overall intention, see Fridland (forthcoming).

⁸ Stanley and Williamson (2001) have argued that knowing how reduces to propositional knowledge. I agree with Ruffitt (2003) that the linguistic evidence that Stanley and Williamson advance for their claim is compatible with the possibility

non-propositional. We can regard having the skill to play the violin as knowing how even if it (in part) involves propositional knowledge. But at least sometimes knowing-how seems to be non-propositional. A baby learns how to grasp a toy, but she does not therefore have propositional knowledge about grasping: she does not have the self-referential belief that she is grasping a toy; I know how to walk, but I don't have propositional knowledge about how to walk.⁹

I will call control when it involves a self-referential belief 'belief-control'. Belief-control is not simply a conjunction of two conditions: belief and control. After all, we can have the belief that we are doing something (e.g. breathing) without it being intentional. We can even have such a belief and control the behaviour, but the control would not be intentional control. A person with some anatomical knowledge may know how to contract a particular muscle (by walking, say). She would then believe that she is contracting the muscle, and control the process. But what she does intentionally in this case is walking. The contraction of the muscle is the intended result, but it isn't an intentional movement.¹⁰ Thus belief-control cannot be the conjunction of belief and control, but it is control (at least in part) in virtue of having the belief. Intentional control is different *qua* control from non-intentional control because it involves the belief. A sleepwalker, for instance, may control his walking quite effectively, but he does not control it intentionally. His control over his movements may not be inferior to intentional control, but it is different. I would like to suggest that

that, sometimes at least, knowing how involves non-propositional knowledge, and that there are reasons to think that it in fact does.

⁹ Stanley (2011) suggests that I may misunderstand what having propositional knowledge involves. Perhaps I cannot *explain* how I walk, but I could at least do it, saying 'this is how you do it'. I know that this is a way of doing it. That suffices for propositional knowledge. However, not even that is true of the baby who nonetheless learns to, and at some point knows how to, grasp a toy. Furthermore, Wiggins (2012) argues persuasively that when we acquire a skill it may result in some propositional knowledge, but we cannot get another person to acquire the skill just by teaching her those items of knowledge. The knowledge that having the skill consists in goes beyond the propositional knowledge that we acquire in developing it.

¹⁰ I am grateful to Anthony Price for the example and for pressing me on this point.

we act intentionally if, and in so far as, we have belief-control over our behaviour. When a person has acquired a skill, she has learned to control part of her behaviour intentionally. But once she has this skill, she may be able to exercise the skill unintentionally (like the sleepwalker).¹¹

Consider the three examples again which defied the earlier attempts at explaining intentional agency. How do they fare under the new criterion?

Doodling: When I doodle, I control what I am doing. Do I belief-control it? That is: do I believe that I doodle (or that I'm drawing a tree, or a face) and control my behaviour in virtue of having this belief? I think the answer is yes. I will say more about this in §iv.

Expressive Actions: When a person acts out of an emotion, she isn't always in control of what she is doing. It is certainly possible to lose control when in the grip of a strong emotion. However, when an agent loses control, her behaviour also ceases to be intentional. Hursthouse's examples are all of agents who are in control of what they do, and they also believe that they are doing those things. Do they control the behaviour in virtue of having the belief (not necessarily the belief that they are expressing anger, but the belief that they are throwing the tin-opener out of the window, say)? Again, I think the answer is yes (and I come back to the point in §iv). Thus, while not all expressions of emotions are intentional, there is expressive *action*, and to it the belief-control criterion applies.

Frankfurt's Spider, and Animal Actions Generally: Are intentional actions performed by animals done in the belief that they are doing what they do? I see no general reason to doubt it. But the question whether other animals (or which other animals) can have self-referential beliefs and control their behaviour in part through those beliefs is an empirical one. I will not pursue it any further here.

Thus the belief-control criterion applies to at least two of the three examples, and perhaps even to all three. The claim I tried to make plausible is that the unifying feature of all intentional

¹¹ I use 'skill' in a very broad sense here applying to highly skilled activities (playing the violin) and more common ones (walking) alike.

agency is that the agent belief-controls it. If the expression of an emotion is beyond the agent's control —as when she trembles, or blushes, say—what she does is not intentional. If an agent controls the precise movement of her fingers (as the surgeon and the violinist do), but doesn't have any belief regarding these movements, the control is not intentional—at least not qua movement of a finger.

Yet if intentional control isn't the only kind of control, what makes it special? What is the philosophical interest in the difference between a person who walks intentionally, believing that she is walking, and the sleepwalker who doesn't have such a belief? I will answer this question in §iv, but want to point out a kinship between this way of understanding intentional agency and that of Anscombe and others.

It has often been claimed that intentional agency involves a certain kind of non-inferential knowledge of what one is doing. Kieran Setiya, inspired by Anscombe, argues that a mark of intentional agency is

belief: When someone is acting intentionally, there must be something he is doing intentionally, not merely trying to do, in the belief that he is doing it (Setiya 2007, p. 26 and *passim*).

As Setiya sees it, **belief** is a necessary truth about intentional agency that any account of it must be able to explain. **belief** is Setiya's version of Anscombe's claim that intentional agency involves practical knowledge, based on the observation that Anscombe's why-question is refused application if an agent when being asked, 'Why are you ϕ -ing?' answers, 'I was not aware that I was ϕ -ing'. As Setiya sees it, it is not actually the lack of awareness, but the agent's lack of belief that she is ϕ -ing which explains why an agent does not ϕ intentionally when she rejects the why-question in this manner.¹²

¹² That is, 'I wasn't aware that I was ϕ -ing' should be understood as indicating that the agent didn't believe that she was ϕ -ing. We can have beliefs that we are not aware of. So even though we express ourselves using the term 'aware' when we want to point out that something we did was not intentional, awareness is really a red herring here.

If intentional agency involves belief-control, this fact (if a fact it is) may help to corroborate Setiya's claim.¹³ It may also help to explain **belief**: because we act intentionally only when we have some degree of control over our behaviour, and because the relevant kind of control is belief-control, **belief** is true.

But isn't there a glaring problem here?¹⁴ There is a certain kind of counterexample to any account of intentional agency that involves belief or knowledge, illustrated by Donald Davidson's carbon copier: 'In writing heavily on this page I may be intending to produce ten legible carbon copies. I do not know, or believe with any confidence, that I am succeeding. But if I am producing ten legible carbon copies, I am certainly doing it intentionally' (1978, p. 92). So Davidson does not need to believe that he is making ten carbon copies in order to make them intentionally. But compare this case with one in which Davidson intends to make seven copies, but cannot be bothered to count the carbon sheets. So he ends up making ten copies. In this case he wouldn't have made the final three copies intentionally, even though everything he does is the same as in the original example. The difference must be in his mental states, and there certainly is a difference in intention. If the view I am expounding here is right, then there must also be a difference in beliefs.¹⁵ In the original case he must believe that he is pressing so very hard on the page that it is at least possible for him to make ten copies; in the revised case, he must believe that he is pressing hard enough to make seven copies.

There remains a question why belief-control is in any way significant. I will venture an answer to this question now.

IV

¹³ However, Setiya's own explanation (2007) of why this is is very different from mine.

¹⁴ I am grateful to Michael Bratman for pressing me on this point.

¹⁵ As will become clearer in §iv, I am not trying to give a reductive account of intentions here: that is, I do not argue for the view, nor do I believe it true, that intentions are beliefs.

Intentional Agency. Human intentional behavior is, at least in principle, subject to reasons.¹⁶ While in two of our three examples (doodling and expressive actions) the agents don't act *for* reasons, they could nonetheless respond to reasons *against* acting in this way, and control their behavior accordingly. When they find that the doodling gets on other people's nerves, or when the angry action is witnessed by a small child who is frightened by it, they can stop. Reason-responsiveness requires belief-control. In order to respond to a reason (that a person believes she has), she must know what she is doing (or at least have a belief about it). If I believe that there is a reason not to undermine my neighbour's ability to concentrate during a philosophy talk, and therefore a reason not to doodle, but don't believe that I am doodling, I couldn't rationally respond to that reason. Intentional behaviour is at least potentially directly responsive to reasons. In that regard, it is different from behaviour which only involves a conjunction of belief and control: the person who believes that she contracts her muscle, and controls it by walking. If there were a reason against contracting the muscle, she could respond to it only indirectly by stopping walking, not by stopping contracting the muscle. Thus my hypothesis here is that control of one's behaviour in virtue of having a belief (controlling it intentionally) is a condition of being able to respond to reasons directly. Take the sleepwalker again: he may well have reason not to walk out into the cold, but since he doesn't believe that he is doing it—since he doesn't do it intentionally—he cannot respond to that reason.

The interest in the subset of those doings where we have belief-control—that is, in (human) intentional agency—is that it introduces the possibility of guiding one's behaviour in the light of reasons.¹⁷

¹⁶ I suspect that at least some other animals can act for reasons, but I am not sure of this. I therefore don't include them here. As a result, the explanation of belief-control in this paragraph may not fully generalize: if belief-control is a condition of intentional animal agency, but other animals don't act for reasons, my explanation of the belief condition may not apply to them.

¹⁷ David Owens suggested to me that there might be another philosophical interest here: understanding responsibility. What an agent is responsible for (what is attributable to her qua agent) are her intentional actions. But while I agree that we

As Michael Bratman has pointed out to me, I am in danger of mixing two 'traditions' of understanding intentional agency: some have argued that knowledge or belief is crucial to intentional agency; others suggest that control or guidance by intention is. This may be so. Sorting out the differences and the arguments on each side will have to wait until another occasion. But I hope that I have shown that we need to go beyond guidance by intention to understand the unity of intentional agency. It may thus turn out that the two traditions are not incompatible after all.

V

Acting with an Intention. Having explained the unity of intentional agency, I will turn from the borderline cases of the three examples, and look at the most developed kind of intentional behaviour — acting with an intention — and at its relation to the reasons for which a person acts.

This kind of intentional agency—a subset of the behaviour over which we have belief-control—has been at the centre of philosophical attention. While agents when they act intentionally, but not with an intention, respond to reasons only in a very rudimentary way (e.g. to reasons to stop), acting with an intention is being guided by reasons. Or so I will now argue. The claim I am going to defend is:

reason: If a person, P , acts with an intention to ϕ , there is (as P believes) a reason to ϕ .

I will begin by distinguishing two senses of acting with an intention, and then ask for each of them how it relates to the reasons for which a person acts. The most common use of the expression 'acting with an intention' is in sentences like 'Susan writes to David with the intention of inviting him to her party'. Here 'with the intention of' can be replaced by 'in order to'. There

need an account of intentional agency to understand responsibility, I do not believe that it determines the scope of responsibility in such a straightforward way.

is something that Susan wants to bring about by writing to David, and this is what she intends in writing to him. Let's call this 'further intention'. There is a less common way of using the expression 'acting with an intention' in sentences like 'Susan writes to David with the intention of writing to him'. Here acting with an intention contrasts with doing something unintentionally, accidentally or by mistake. Call this 'basic intention'.

Further Intentions. If we know which further intention an agent pursues in her action we also know something about her reasons (the reasons she believes she has) for acting in this way. If we were to ask Susan, 'Why do you write to David', she would answer, 'To invite him to my party', thereby giving us her reason for writing to him. So if there is a further intention with which someone acts there is also (as she believes) a reason for her acting as she does.

Basic Intentions. How about basic intentions? Must an agent when she acts with a basic intention believe that there is a reason for doing so? This is not as obvious as with further intentions, since the description of the basic intention does not bring out what the agent's reason for acting is. However, one feature of acting with a basic intention is that there is something that the agent regards as success or failure of her action: if Susan, intending to write to David, accidentally sends her message to Peter, she would have failed to do what she intended to do. That Susan regards this as a failure indicates that she sees a reason for writing to David, and not to Peter.

Often the reason for the action that is done with a basic intention is that it facilitates achieving an end. Thus there is also a further intention. The reason why Susan intends (basic intention) to write to David rather than to Peter is that she intends (further intention) to invite David (but not Peter) to her party. In those cases, since the agent believes she has a reason to achieve the goal, she may also believe that she has a reason to write as a facilitative step towards achieving her goal. (Whether or not she is right about the latter depends, among other things, on whether she is right about the former.)

But let's consider an action done with a basic intention where there is no further intention: passing the time by whistling a tune, intending to whistle that tune, when there is nothing that counts as success or failure. Alternatively, I may whistle because I enjoy doing it, but only in so far as I get it right. That I enjoy it gives me a reason to continue (just as the fact that you don't enjoy it gives me reason to stop), and it explains why it matters to me whether I get it right: my enjoying the activity depends on it.¹⁸ In the first scenario the whistling isn't done for a reason (it is like doodling); but in the second there is a reason: the enjoyment which imposes a standard. The differences between the two scenarios may not be sharply carved out, but we should distinguish them nonetheless: when there is a possibility of failure, the agent may act for a reason even when there is no further intention. When there is nothing that would count as failure, the agent is not acting for a reason. Failure and success bring in a normative dimension, and therefore indicate that the agent has a reason for acting as she does, or at least believes she has one. If an agent believes that she shouldn't fail in a particular pursuit, she must see a point in doing what she does —and that may suffice to show that she believes she has a reason for her action.

There is a real distinction here, but it probably doesn't map perfectly on the ordinary use of the expression 'acting with an intention', granted anyway that we rarely use this expression except to refer to further intentions. I will use 'basic intention' to refer to only those cases where the agent does regard not succeeding in acting on the intention as a failure, either because she sees some value in acting in this way, or because she believes that acting in this way facilitates realizing a further intention. In those cases she acts for (what she believes is) a reason. By contrast, when it doesn't matter whether the agent succeeds or fails, she is not acting for a reason, but the behaviour is nonetheless intentional in the same way in which, for instance, doodling is. It does involve belief-control, but no

¹⁸ I assume here that I don't need to have a further intention to enjoy myself (otherwise the example would revert to basic intention when there also is a further intention), but that it is enough that there is a reason to enjoy myself, and that the enjoyment depends on getting it right.

reasons to act. Thus, acting with a basic intention is different from merely having belief-control, because it introduces a normative dimension that belief-control lacks. (There is no failure in doodling, or in expressions of emotion.¹⁹)

If we understand acting with a basic intention (even when there is no further intention) as an action where the agent holds herself to a normative standard of success, then when an agent acts with an intention (further or basic), she acts for (what she believes is) a reason. Thus, with a little bit of terminological tweaking, **reason** is true.

VI

Intentions and Reasons. According to **reason**, when a person acts with an intention she acts for (what she believes is) a reason. How about the reverse: does an agent act with an intention when she acts for a reason? Might it be possible to act for a reason, but not with an intention to do so? Answering this question will shed some further light on the role of forming intentions.

There are always many reasons for and against acting in a certain way. An agent may be aware of at least some of these — we are often

aware of a multitude of reasons for various options when we deliberate. In deliberation, we attempt to figure out which of those reasons to act on. By forming an intention, the agent resolves which reason she is going to act on.

That she resolves the question does not entail that she will act in this way, or that she must believe that she will do so. After all, she can revise her intention later, or forget about it, when the time to act on it comes. Or she can fail to keep track of time without forgetting. But at least for now she has settled which of

¹⁹ Or rather, when there is failure in, say, expressing one's emotions — as when I, rather ineptly, express my gratitude for someone's help — then this is not the kind of expressive action I described above, but a case of acting with the intention of expressing one's emotions. In that case, I do act for a reason, and am therefore subject to standards of success and failure.

her reasons to act for. *If an agent doesn't change her mind, doesn't forget about the intention and keeps track of time, she will at least try to do what she intends to do.*²⁰ Call this conditional claim **resolve**.

Some philosophers (e.g. Harman 1976) claim that an agent who has an intention must believe that she will act as she intends to do. But since we know that changing our mind, forgetting, or failing to keep track of time are live options for each of us, it would, for the most part, be irrational to have such a belief. Rationally, an agent can only believe the conditional claim, **resolve**.

Forming an intention may help one to act in accordance with one's reasons in the way **resolve** indicates: by settling the deliberative question, and by moving closer to action. When the agent does act on her intention, the intention will guide her through the whole process of acting. It doesn't just trigger the action, but guides it in all its stages, including modifications in the light of changing circumstances, or even giving up when realizing the intention becomes too costly. Acting with an intention is the psychological side of being guided by a reason.²¹ This second feature of acting with an intention is the *guidance* that the intention affords. Thus forming an intention has two main effects: it resolves what to do and it guides the agent through a course of action.

When we act for reasons as a result of deliberation we often act with **resolve** and *guidance*: reasons don't resolve among themselves
what a person will do, and while they can guide her, they must do so with the help of psychological states whose function is to

²⁰ As Julia Driver pointed out to me, **resolve** holds only for those intentions where not acting within a certain time frame would count as failing to act as intended. An agent may intend to do something at some yet to be specified time in the future. She may intend to take her daughter on a holiday, say. Intentions of this kind can be fully specified and detailed in certain regards: she may have a plan, except that the time isn't fixed yet. In this case, the conditional would not hold: even if she doesn't change her mind, doesn't forget, and doesn't fail to keep track of time, it doesn't follow that she will act on the intention. But there also isn't any time (within the lifetime of both daughter and mother) where she would have failed to act as intended.

²¹ Thus intentions are not themselves reasons.

monitor the action's progress. This is what forming an intention achieves: it is a way of setting oneself to act in accordance with one's reasons, and, if one acts on the intention, monitoring the action's progress in the light of the reasons.

But forming an intention not the only way of achieving this. Take cases in which the intention would undermine one's ability to do what one has reason to do: when the reason is a reason to act spontaneously, say.²² In such a case, the outcome of deliberation may be that I ought to do something spontaneously—but I cannot achieve this by forming an intention to do so. Or take the violin player: she has a reason to play each and every note of the concerto, and she may go through the sheet music being aware of the succession of the notes. But it would be counter-productive and confusing to form intentions to play each of them. Thus while forming an intention upon deliberation in many cases facilitates acting for the reasons one believes to have, it doesn't always do that. While intentions can be useful in so far as they provide resolve and guidance, they can also undermine the efficiency of the agent (the violinist), or even the whole point of the action (the spontaneous action). We can certainly act in accordance with reasons without intending to do so, or while intending something quite different. And it seems that we can even act for a reason as a result of deliberation without forming an intention to do so—assuming that it is possible to act spontaneously for the reason that one should act spontaneously.

To sum up, I have shown that acting intentionally is a unified phenomenon, and that it is crucial to our ability as rational agents to act for reasons. Yet this is not because acting intentionally *is* acting for a reason. A second explanation of the unified phenomenon also turned out to be false: that an action is intentional if and because the agent acts with an intention. There is however a unifying core: intentional agency involves a special kind of control, the control that an agent has when she acts in the belief that she is doing what she does. Belief-control of this kind opens up the possibility of responding to reasons.

²² I am grateful to Grant Rozeboom for alerting me to examples of this kind.

Furthermore, when an action is not just intentional in this broad sense, but done with an intention, it is done for a reason. Thus while the view that acting intentionally is acting for a reason is false, it is true that acting with an intention is acting for a reason.

I have explored how acting intentionally relates to acting for reasons and to acting with an intention. The reasons in question are reasons to act. We found that when a person intends to do something, it is because she believes she has a reason to do it. Thus, as she sees it, there is something to be said for the action she intends to perform. Recall **reason**: if a person, P , intends to ϕ , there is (as P sees it) a reason to ϕ . We form intentions on the basis of reasons to act (that we believe we have). The currently prevalent view of the relation of reasons to act and intentions may seem at first blush congenial to, and perhaps even providing an explanation for, **reason**. I will now, in the final paragraphs, put the suggested view of intentional agency to work, and show that this is not so, and that the prevalent view is false.

VII

The Symmetry View. Do we act for a reason when we act with an intention, because the very same reason—the reason for acting—is also a reason to form the intention? Do we rationally form and revise intentions in response to our reasons to act? According to a dominant strand in the current discussion the answer is yes: a reason to ϕ is *ipso facto* a reason to intend to ϕ . And therefore, when a person believes she has a reason to ϕ , she would rationally form an intention to ϕ . The proponents of this view argue that reasons to intend should be understood in parallel with reasons to believe. Special cases apart, we form beliefs in response to evidence for the truth of a proposition, and we discard beliefs when there is evidence for their falsity. In parallel, so the view, intentions are formed in response to reasons to act: reasons to act govern the formation of intentions and determine their rationality in much the same way in which evidence governs the formation of beliefs. Call this the *symmetry view*.

It is defended explicitly by a number of philosophers (Hieronymi 2005; Shah 2008; McHugh 2012), but one of its claims is at least implicitly accepted by nearly everyone in the debate: that reasons to act are reasons to intend, and that truth-related considerations are reasons to believe. The dispute is mainly about the question whether those are the only reasons for the attitudes in question, or whether some other considerations provide reasons too (Pink 1991; Schroeder 2012).

The symmetry view aspires to giving a unified account of theoretical and practical normativity based on the commonalities between theoretical and practical reasons. It rests crucially on *one observation*: we cannot form beliefs exclusively for a certain kind of reason, namely, having a belief because it would be good to have it. I cannot form the belief that Cameron is a good prime minister just because it would help me to get a job if I had this belief. I cannot form the belief directly for *that* reason. (A whole variety of demon-induced reasons are of this kind.) In general, I cannot form the belief that *p* just because (as I see it) it would be good if I had the belief that *p* in the absence of any evidence that *p* is true.

Similarly, I cannot form an intention exclusively for reasons that bear only on the value of having the intention: reasons to form the intention to ϕ in the absence of any reason to ϕ . It may be good to have an intention, even when there is nothing to be said for acting as intended. We find examples of this kind in the neighbourhood of the toxin puzzle.²³

The proponents of the symmetry view set out to explain why it is that we cannot form beliefs and intentions directly for reasons that bear only on the value of the attitude. They aspire to providing an explanation which also establishes their main

²³ The toxin puzzle (Kavka 1983): an eccentric billionaire would transfer a million pounds into your bank account at midnight today, if you now intend to drink a (mild) toxin tomorrow. The toxin will not kill you, but it will cause you some discomfort. The billionaire does not require that you drink the toxin, but only that you form the intention to do so. You know now that tomorrow there will be no reason for you to drink the toxin, since either the money is already in your bank account or you won't receive it. Thus you have a reason to form the intention to drink the toxin, but no reason to actually drink it.

claim: that reasons to act are reasons to intend, and possibly even the only reasons to intend.

According to one of the common explanations, normativism, the attitudes are subject to constitutive standards of correctness. In the case of belief, the standard is truth; in the case of intentions, the standard has been dubbed 'to-be-doneness' (Shah 2008). An attitude would not be a belief, so the idea goes, if it didn't respond to evidence. And an attitude would not be an intention if it didn't respond to considerations that bear on the 'to-be-doneness' of actions. Since the reasons that are provided by the value of the attitude itself are by definition independent of truth or 'to-be-doneness', the attitudes do not respond to them: we cannot form the respective attitudes for those reasons.

An alternative explanation is a teleological one: beliefs 'aim' at truth (or knowledge), and intentions 'aim' at actions which are to be done (McHugh 2012, 2013). Both normativism and teleology promise to explain why we cannot form intentions in response to reasons of the toxin puzzle variety. I doubt that either view delivers what it promises, but I will not pursue this question here. Instead I will argue that the main claim of the symmetry view is false, not just that its justifications are problematic (at least with regard to reasons to intend).

VIII

The Facilitative Reason to Intend. As I argued in §vi, forming intentions has characteristic benefits: it resolves which reason to act for out of the many reasons we often have for or against acting in a certain way. It thereby settles what the agent will do, unless she changes her mind, forgets, or fails to keep track of time. And, assuming the agent does act on the intention, it guides her through the course of action. Therefore, when the agent has a sufficient or conclusive reason for doing something, and doing it will be more secure if she forms an intention, then she has a reason to form this intention. It facilitates acting in accordance with one's reasons. This is a kind of instrumental reason, since

the reason is provided by the psychological benefits of forming the intention: it can be useful to have it. As with instrumental reasons generally, there is such a reason only when there is at least a sufficient reason to act as one intends to do.²⁴ But the facilitative reason does not derive from the reason to act in the way the symmetry view claims reasons to intend derive from reasons to act. It isn't simply the reason to act that provides the reason to intend, but the benefit of forming an intention when there is a reason to act as intended.

Let me make this clearer by focusing on an example. Let's assume that Julie has reason to write a novel. It is hard to imagine that anyone could write a novel without intending to do so. If Julie has a sufficient reason to write (or to try to write) a novel, she also has a facilitative reason to intend to do so. The facilitative reason is a reason to form a basic intention: an intention to write a novel. But the reasons for writing the novel are different. There is no reason to write a novel because it is writing a novel. The reasons for writing might be that Julie is a gifted storyteller, that she has seized on an important topic, that she needs to earn a living, or that it would develop her talents. Thus her basic intention (to write a novel) does not mirror the reasons for writing the novel. But presumably she would not only have the basic intention to write, but some further intention in writing her novel as well. Would that derive directly from the reasons to write? It wouldn't, since she needn't have any particular further intention: she would comply with her reason to develop her talents even if she were writing with the intention of telling the story. She needn't intend to develop her talent even if this is a sufficient reason for her to write. If she were to form the intention to write the novel to tell an important story, she would not have failed to respond to her reasons. The reason to intend does not then derive from the reason she has to write because it would develop her talent. But according to the symmetry view, a reason to intend to ϕ derives from any sufficient reason to ϕ .

Furthermore, the facilitative reason is of the wrong kind for the

²⁴ That is, I presuppose here that an agent, *A*, has an instrumental reason to ϕ in order to ψ , when ϕ -ing is sufficient for ψ -ing, iff she has at least a sufficient reason to ψ .

proponent of the symmetry view: being an instrumental reason that the author has because of the psychological benefits of intending (re- solve and guidance), it is of a kind with the reasons that the symmetry view rejects. It is based on a value of intending, which is independent of the value of acting as intended. Thus what is perhaps the most pervasive reason to intend is completely at odds with the symmetry view.

IX

Counterexamples to the Symmetry View. The symmetry view is mistaken about reasons to intend, as the novel-writing example brings out: complying with one's reasons does not require that one should do so intentionally—at least not in general.

This is a plain rejection of the symmetry view, but perhaps also the most obvious objection: the symmetry view is false because there are clear counterexamples. Here are three of those.²⁵

- (1) Many reasons to act are neither explicitly nor implicitly reasons to act intentionally (as the novel-writing example in §viii has shown).
- (2) Some reasons to act cannot be complied with by forming an intention. Doing so would be self-undermining (as the earlier example in §vi of acting spontaneously has shown).
- (3) Reasons not to do something —reasons for omissions — are not reasons to intend not to do it. There is no reason to intend not to kill or to betray, say. If it never crosses your mind whether or not to kill, you haven't failed to respond to your reasons.

From these counterexamples we can conclude right away that it is not true that we have a reason to intend to ϕ when we have a

²⁵ For a more detailed discussion see Heuer (forthcoming).

(sufficient) reason to ϕ . When is there a reason to intend to ϕ then?

X

Reasons to Intend. The novel-writing example has shown that there isn't a reason (other than the facilitative one) to intend to ϕ when there is a reason to ϕ , as long as the reason to ϕ isn't a reason to ϕ intentionally. But there are some actions which we *cannot* perform except intentionally: giving a gift, thanking or marrying someone, say. If I give you something which I do not intend as a gift, then it isn't. Thanking a person is thanking her intentionally. If there is a reason to perform an action of this kind, it is *ipso facto* a reason to do so intentionally.

There are also some actions that we *shouldn't* perform unless we do so with a particular intention: a reason for supporting a vulnerable person may require that one doesn't act for personal gain; there may be a reason to visit a sick friend only when the intention is to cheer her up, but not, for example, to win a bet. Thus we get the following schema:

i. *Reasons for future-directed intentions*

A person, P , has a reason to intend to ϕ if

(1) P has a (conclusive or sufficient) reason to ϕ , and in-

ii. *Reasons for intentions in action*

A person, P , has a reason to intend to ϕ if

(2) P has a (conclusive or sufficient) reason to ϕ intentionally

(for example, a reason to give someone a gift).

(3) P has a (conclusive or sufficient) reason to ϕ , provided it is done with a particular intention, or for certain reasons, but not others (for example, visiting a sick friend to cheer her up, but not to win a bet).

The symmetry view does not account for any of those reasons, and postulates reasons that we don't have. It should therefore be rejected.

XI

Conclusion. On the prevalent view of the relation of intentions and reasons to act, we have reason to intend whenever we have a sufficient reason to act. I have shown that this is false because there are clear counterexamples of reasons to act which do not provide reasons to intend, as well as of reasons to intend which do not simply derive from having reasons to act.

What's more, the symmetry view misunderstands the relation of practical reasons and intentions. It conceives of it as a normative relation: we *should* form an intention when we have a sufficient reason to act. In §§iv–vi, I have shown that we cannot act for reasons except by acting intentionally. Acting intentionally is our main way of conducting ourselves, and this is so because, as rational agents, we act (much of the time) for reasons. The kind of control that comes with intentional agency allows us to respond to reasons. When we act for reasons we act intentionally, not — special cases apart — because we have a reason to act with an intention (intention in action), but because we cannot act for reasons in any other way. The aspiring author cannot write a novel without intending to do so. Sometimes we form intentions for reasons — in particular, we form future-directed intentions for facilitative reasons — but the symmetry view does not account for those.

The symmetry view also ignores that acting for a reason is only one way of complying with our reasons. Often there is nothing amiss as long as we act in accordance with our reasons (e.g. reasons for omissions). In those cases we can comply with our reason to act even when we have no particular intention to do so.

The symmetry view is mistaken all the way: it stipulates reasons to intend that we don't have, and it does not account for the reasons to intend that we do have.

Coda. Reasons to intend because there is value in having the intention (independently of the value of what is intended) seem straight-forward reasons to form an intention. Why is it that we sometimes can't follow them, at least not directly? (We can follow them when there is a reason to act as intended as well, as the facilitative reason shows.²⁶) Here is an explanation that seems plausible to me: the reason for forming an intention in the toxin puzzle is a reason to form a future-directed intention. As I explained in §vi, future-directed intentions will become intentions in action, provided the agent doesn't change her mind, doesn't forget, and keeps time. So the future-directed intention to drink the toxin in the toxin puzzle would become the intention in action when the time for drinking arrives, unless the agent changes her mind, or forgets, or fails to keep track of time. If it were to become her intention in action, she would act with an intention, but for no reason whatsoever. After all, when the time of drinking comes there is no reason to drink the toxin. If what I said in §v is correct — that is, if **reason** is true — then it is not possible to act with an intention but for no reason. So she may instead change her mind, and decide not to drink the toxin once the time comes. But she wouldn't change her mind in response to new information: all the information is in (we assume) at the time when she forms the future-directed intention. So she would change her mind for no reason. That may well be possible. But in this particular case, if the agent is rational, she would know ahead of time that she will change her mind for no reason. That seems tantamount to intending to do something that you know you won't do. And that may not be possible.²⁷

²⁶ We can also follow them when we can acquire the intention by acting in a certain way — acting can be a means of complying with a reason to intend: if you offer me £10 if I intend to raise my hand, I can simply raise my hand. Since I do so intentionally I would win the money. See Hieronymi (2005).

²⁷ I presented earlier versions of this paper at a conference on the moral significance of intentions at the University of Leeds, and at a workshop on rationality, agency and morality at the University of Kent at Canterbury. I am grateful to audiences there and at the meeting of the Aristotelian Society for their insightful questions and comments. I would also like to thank Dorothea Debus and Joseph Raz for very helpful comments and conversations. The research on this paper was funded by the

Philosophy
University of Leeds
Woodhouse Lane
Leeds LS2 9JT
UK

References

- Anscombe, G. E. M. 1963: *Intention*. 2nd edn. Oxford: Basil Blackwell.
- Bratman, Michael 1987: *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- Davidson, Donald 1963: 'Actions, Reasons and Causes'. Reprinted in Davidson 1980.
- 1978: 'Intending'. Reprinted in Davidson 1980.
- 1980: *Essays on Actions and Events*. Oxford: Oxford University Press.
- Frankfurt, Harry G. 1978: 'The Problem of Action'. *American Philosophical Quarterly*, 15 (2), pp. 157–62. Reprinted in his *The Importance of What We Care About: Philosophical Essays*, pp. 69–79. Cambridge: Cambridge University Press, 1988.
- Fridland, Ellen forthcoming: 'They've Lost Control'. To appear in *Synthese*.
- Harman, Gilbert 1976: 'Practical Reasoning'. *Review of Metaphysics*, 29 (3), pp. 431–63.
- Heuer, Ulrike forthcoming: 'Reasons to Intend'. To appear in Daniel Star (ed.), *The Oxford Handbook of Reasons and Normativity*. Oxford: Oxford University Press.
- Hieronymi, Pamela 2005: 'The Wrong Kind of Reason'. *Journal of Philosophy*, 102 (9), pp. 437–57.
- Hornsby, Jennifer 1993: 'On What's Intentionally Done'. In Stephen Shute, John Gardner and Jeremy Horder (eds.), *Action and Value in Criminal Law*. Oxford: Clarendon Press.
- Hursthouse, Rosalind 1991: 'Arational Actions'. *Journal of Philosophy*, 88 (2), pp. 57–68.
- Kavka, Gregory 1983: 'The Toxin Puzzle'. *Analysis*, 43, pp. 33–6.

British Academy, and by the ERC advanced project 'Distortions of Normativity'.

- McHugh, Conor 2012: 'Control of Belief and Intention'. *Thought*, 1 (4), pp. 337–46.
- 2013: 'The Illusion of Exclusivity'. *European Journal of Philosophy*, 22 (3). doi: 10.1111/ejop.12032.
- Pink, T. L. M. 1991: 'Purposive Intending'. *Mind*, 100, pp. 343–59.
- Raz, Joseph 2011: *From Normativity to Responsibility*. Oxford: Oxford University Press.
- Rumfitt, Ian 2003: 'Savoir Faire'. *Journal of Philosophy*, 100 (3), pp. 158–66.
- Schroeder, Mark 2012: 'The Ubiquity of State-Given Reasons'. *Ethics*, 122 (3), pp. 457–88.
- Setiya, Kieran 2004: 'Explaining Action'. *Philosophical Review*, 112, pp. 339–93.
- 2007: *Reasons Without Rationalism*. Oxford: Oxford University Press.
- Shah, Nishi 2008: 'How Action Governs Intention'. *Philosophers' Imprint*, 8(5).
- Stanley, Jason 2011: 'Knowing (How)'. *Noûs*, 45 (2), pp. 207–38.
- and Timothy Williamson 2001: 'Knowing How'. *Journal of Philosophy*, 98 (8), pp. 411–44.
- Wiggins, David 2012: 'Practical Knowledge: Knowing How To and Knowing That'. *Mind*, 121, pp. 97–130.

