

4

Entitlement to Reasons for Action

Abraham Seshu Roth

Intentional action traditionally is understood in terms of reasons explanation.¹ Crudely put, when I ϕ intentionally, I ϕ for a reason, where the reason explains the point of ϕ -ing. I contend that on occasion one's *entitlement* to certain reasons had by someone else allows for those reasons to figure in the intentional description of what one does. The entitlement secures for those reasons an explanatory role they otherwise would lack. Normally, the reasons for which I act are *my* reasons; I represent goal states and the means to attaining them, and these reasons guide me in action. If I haven't yet taken up your reason and made it mine by representing it for myself, then it may seem mysterious how it explains what I do. Nevertheless, I argue that sometimes what one does should be explained this way, and I take some steps toward showing how this is possible.

I'm concerned here with *explanatory* reasons, and this makes defending the proposal more challenging. Merely *justificatory* reasons might be given to put an agent's ϕ -ing in a good light, but they don't purport to explain why the agent did what she did or describe what she was up to in doing it. You may have all sorts of reasons for thinking that my ϕ -ing is a good thing. No one would find it problematic for you to cite them as some justification or defense for what I've done, even if I were not aware of them. But if I don't have those reasons in mind, then it is controversial to cite them to *explain* my ϕ -ing, to make sense of what I was up to when I ϕ -d.²

¹ Davidson (1980), Anscombe (1963).

² The relevant notion of explanation is the sort of rationalizing explanation that philosophers such as Davidson and Anscombe have tied to intentional action. This underlies the novelty and significance of the central thesis of the chapter. There are, of course, other kinds of explanation of action that might be offered up by cognitive or social psychology, and it might not be news that the agent needn't represent for herself explanations of her actions so understood. It's a good question just how those explanations are related to the reasons explanations that are my concern here. I won't try to answer that interesting

It is not surprising, then, that in discussions of the explanation of action, we do not seem to encounter the sort of possibility to be defended here. Davidson, for example, insists that what explains the action is the agent's primary reason, which shows "what the agent sees in the action". He understands the primary reason to consist of the agent's desire (or some other pro-attitude) regarding actions with a certain property, along with a belief that the action has this property (Davidson 1963: 687). But I am suggesting that the reason that can sometimes be invoked in explanation is not something that the agent sees in the action, and it is not a reason that the agent possesses in Davidson's sense. Rather, it is a reason had by someone *else*.

Anscombe likewise doesn't seem to consider the possibility.³ She describes a case where a host tells a guest to hang a hat on a hook. She finds that it is often unclear in these sorts of cases whether what we have is a mere automatic response to be explained by causation, or a full-blown intentional action she thinks is to be explained by reasons and not causes. But she says that if one were forced to decide, "the more the action is described as a mere response, the more inclined one would be to use the word 'cause'; while the more it is described as a response to something as having a *significance that is dwelt on by the agent, or as a response surrounded with thoughts and questions*, the more inclined one would be to use the word 'reason'."⁴ My concern here is not whether reasons explanation is or is not causal explanation. What is of interest is that Anscombe uses language here that is suggestive of the agent having in mind the reason for the action. Thus, for Anscombe too, it seems the agent must in some sense possess the reason in order for it to explain the action.

My proposal, in contrast, is that sometimes one may act for reasons that one doesn't possess. *Entitlement* to another's reason allows that reason to figure in the explanation of some of the things that one does. The natural home for this idea of entitlement would be in a theory of shared agency. That's not to say that the idea of entitlement has no precedents in our understanding of individual agency. Indeed, the strategy for clarifying the proposal and making the case for it will be to bring to the fore aspects of the interplay between reasons and intentions in the individual case that have not received the attention they deserve. This is what I go on to do in the next section.

question on this occasion. I take it that what I'm saying here might nevertheless be of interest for anyone except perhaps the crudest advocate of a simple-minded eliminativism about reasons explanation.

³ She does consider the possibility of the agent being given a reason and accepting it (Anscombe 1957: 330–1).

⁴ Anscombe (1957: 331, emphasis added). As she makes clear shortly thereafter, the thoughts in question are ones that contribute to answering the appropriate why-question—the one that seeks reasons for the action.

I then consider the interpersonal case in section 2. The scope of entitlement is not unlimited, and in the final section I explore a couple of ways in which we might understand how it is constrained.

4.1 ACTING ON A PRIOR INTENTION

As a preliminary, consider what might be called *guidance* cases, where one attends to the progress one is making toward some goal and guides what one is doing in light of this attention.⁵ For example, I am in a seminar and want to ask a question. I raise my hand in such a way as to get the speaker's attention. How high I raise my arm, my timing in executing it, and whether it is supplemented perhaps acoustically or by a wave of the hand or a shift in posture, will all depend on the progress I'm making toward getting her attention. If someone were to ask why I am raising my hand, I might invoke the 'in order to' locution and say something like this:

- (1) I'm raising my hand in order to get the speaker's attention (in order that she may call on me).

Here, we are supposing that a consideration that explains why one is acting—the reason for which one is doing whatever it is that one is doing—is explicitly and actively represented in the agent's action plan. The reason amounts to the goal or aim of action, and the agent (or causal mechanism within the agent) guides movements or sub-actions to bring about the explicitly and actively represented aim. The **attentive guidance model** of reasons explanation holds that *any* reasons explanation of action must conform to this picture of guidance.

But this guidance model goes too far. Sometimes a consideration that might serve as an explanatory reason doesn't seem to be explicitly and actively represented. In this sense, the agent doesn't have in mind the further reason when performing the action. Different sorts of cases are relevant here,⁶ but let me illustrate with acting on a *prior* intention, which will occupy most of my discussion. Suppose that I'm scheduled to get a new fridge installed and, to make way for it, I must move the sofa to the edge of the living room. So, I decide I will do this later when I get a chance. When the time comes and I'm moving the sofa, do I necessarily have in mind that

⁵ Frankfurt (1978), Wilson (1989). This is not to say that these authors advocate the guidance model discussed below. For some discussion of why guidance is not just a matter of pure feedback, see Jeannerod (2006).

⁶ Action with ends that are quite removed; policies; habitual action (where one attends neither to the performance itself nor to some more removed end).

I'm doing this in order to make way for the fridge? Yes, most likely. But do I *have* to? If I can do the job well enough without consciously attending to my progress toward the goal of making way for the fridge, then I don't see why this is necessary. It is cognitively burdensome to have to keep track of all the reasons I may have for whatever it is that I'm doing (Bratman 1987, Harman 1986). Often I just want to focus on the act; but sometimes I like to think about or daydream about other things while performing the act. Either way, the further goals of the act need not occupy my attention. So let us imagine that I'm not thinking about making way for the fridge as I push the sofa to the edge of the room. Nevertheless, if someone were to ask why I was doing this, it seems that the appropriate reasons explanation would be something like:

- (2) I am moving the sofa in order to make way for the fridge.

It would be appropriate to describe *what I'm doing* as making way for the fridge. And I maintain this even though I don't explicitly and actively represent this reason at the time of action.

For another case, let's turn to an example from what we might call middle-aged philosophy of action, and consider a case of forgetting what one is doing. Suppose that I walk down the hall from my office to the main department office to retrieve a printout. By the time I reach the main office, I have a senior moment and do not remember what I've come for. I end up peering into my empty mailbox and having a conversation with a colleague instead. Nevertheless, the explanation for my walking down the hall would be something like:

- (3) I am walking down the hall in order to retrieve the printout

even though if someone were to ask me I might already have forgotten and be unable to give this answer. This case reinforces the lesson drawn from acting on prior intentions more generally, as in the sofa example. Namely, the scope of reasons explanation extends beyond what is entailed by the attentive guidance model.⁷

Resisting the intrapersonal (diachronic) entitlement claim: Are we being too quick to dismiss the attentive guidance model? Can we, for example, re-interpret the sofa case so that it conforms to that model? That is, assuming that when moving the sofa I'm not explicitly and actively representing the goal of getting the fridge through, might we point to something else that is explicitly represented and serves to guide my performance? It might

⁷ Many instances of negative action or omission also serve as counterexamples to the attentive guidance model. My not going to a party can be intentional; my reason is that I don't want to see X there. However, I don't have to attend to or guide what I do to avoid going to the party. I just stay at home as I normally do.

be that when I move the sofa, I'm also thinking about creating more flow or a more pleasing living space. But we can just stipulate that in this case I don't have such other reasons for moving the sofa. Still, we might try to re-interpret the case by appealing to the fact that, although I don't have in mind that I'm making way for the fridge, *I am acting on the prior intention* to move the sofa. Acting on a prior intention is an inescapable element of the sort of case I'm describing, and so I can't just stipulate it away. So perhaps we have the following reasons explanation:

- (4) I'm moving the sofa in order to act on my prior intention (to move the sofa).

There may be cases like this. Perhaps I tend to be lazy or weak-willed, and am making an effort to take my resolutions and decisions more seriously. In this context, given that I have decided and now intend to move the sofa, I should stick with it. (Perhaps my therapist tells me that I should practice following through on my intentions...) That *might* be a way of making sense of what I'm doing here. But the peculiarity of such a rationalization suggests that this strategy of re-interpreting the case won't generalize.

Part of what makes the proposed re-interpretation problematic as a more general approach is that it distorts the subject matter or justificatory structure of the reasons. In the case at hand, I'm moving the sofa in order to make way for the fridge. *That's* the reason. The justificatory structure faces outward toward the worldly concerns of home improvement; it's not (normally) a matter of self-improvement—that is, it's not about my resoluteness or remedying the lack thereof.

In this regard, we should keep in mind what it is that we are normally interested in when we ask why someone is φ -ing: we're asking what the point of φ -ing is. So we would be frustrated if the answer were something like, he φ -d in order to fulfill his earlier decision or prior intention to φ . We would immediately ask why he had so decided. At this point the advocate of the attentive guidance model might acknowledge that the consideration about making way for the fridge figured in my prior decision to move the sofa, but insist that by the time I actually move the sofa, if I am no longer thinking of the fridge, then the only *explanatory* reason on the scene is that concerning fulfillment of the prior intention (to move the sofa). But I think it would be odd for the reasons for intending/decision to come apart so easily from the reasons for the corresponding action.⁸

So, my response to the re-interpretation of the example is that it's committed to a mistaken inward- or self-directedness of the reasons for the

⁸ That's not to deny that they sometimes do come apart, as the toxin case illustrates (Kavka 1983).

action.⁹ But if that's the charge, can the problem be ameliorated with the suggestion that fulfilling a prior intention to φ counts as a reason for φ -ing on the grounds that fulfilling prior intentions generally promotes better outcomes? Thus the explanation might be something like:

- (5) I'm moving the sofa in order to act on my prior intention to move it (and I'm acting on my prior intention because doing so tends reliably towards better outcomes).

That *may* be a reason for acting, but again it strikes me as a rather odd one. In the normal course of things, the general reliability of the policy of acting on prior intentions (even if true) does not figure in the rationalization of a particular action. It is, rather, taken for granted as a presupposition for reasons explanation of action, and not normally what one is looking for when asking why the agent is doing something. This explanation might be given by self-conscious action theorists. I might use it as a justification if I'm being particularly reflective and wondering whether I've been pursuing my ends in an effective or efficient manner. I suspect, however, that this is not what most agents are concerned with, and it isn't what I'm up to when for example I move the sofa.

I conclude that the natural context in which reasons explanations are requested and given shows that it is implausible to appeal to something like *implementing an intention* as a reason for the action. That suggests that we should interpret the cases we discussed as we did initially: as ones for which the attentive guidance model does not apply. Not all explanatory reasons are explicitly and actively represented by the agent in the performance of the action.

Neither of the proposed re-interpretations of the sofa case is plausible. Still, we might wonder what positive consideration might be given for thinking that considerations that don't figure in attentive guidance may nevertheless serve as explanatory reasons. Here's one thought. Though I don't have in mind at the time of action that the reason for moving the sofa is to make way for the fridge, that reason nevertheless serves as a standard for whether I am successful in performing the action. And it's not just any standard. There are, of course, many standards against which the action might be judged. It might be executed gracefully enough to meet the standards for providing aesthetic pleasure for those witnessing my efforts. Or, it might meet the standard of reassuring housemates that I am willing to do my share of what needs to be done around the house. But these weren't

⁹ A further worry becomes evident if we recall that the proposed re-interpretation is tied to defending the attentive guidance model. The proposal holds, implausibly, that the agent necessarily is attending to the prior decision or intention.

(at least in this case) the standards relevant for defining what *I* was up to, and so standards that *explain* what I'm doing. The relevant standards define whether I succeed at what I'm doing and are those that at least in principle play a *regulative* role for the agent. Although I may not have standards in mind when I perform the sofa moving, I may afterwards assess whether I for example moved the sofa far enough so as to allow the fridge to get through. If not, I may have to have another go at it.¹⁰

So, again, the normal structure of the reason explanation when I'm acting on the prior intention was:

- (2) I'm moving the sofa in order to make way for the fridge.

Though the consideration of making way for the fridge is not in this instance implicated in the active guidance of my sofa-moving, I am *entitled* to it as part of the intentional description of what I'm doing in moving the sofa, and as a reason for why I'm moving it. Entitlement relates one to reasons in a way to be distinguished from the active guidance that we have in the attentive guidance model; it does not require explicit representation at the time of action.

The line of thought here starts with a claim about the structure or subject matter of explanatory reasons in the case where the agent is acting on a prior intention. (We might put this point by saying that the intentionality of the action is richer than that of merely acting on some prior intention.) If a reason explains an action, then the agent performs that action for that reason. If the agent acts for that reason, then either the reason figures in the attentive guidance of the action (i.e. the reason is explicitly and actively represented), or else the agent is related to the reason in some other way. Call this other way of relating to reasons *entitlement*. Sometimes an agent acts for a reason, but does not do so via attentive guidance. So, sometimes an agent has an entitlement to a reason that explains his or her action.

The argument is meant to show *that* there must be some such entitlement. I am not offering a theory or account of what this entitlement consists in or how it's possible to have this entitlement. Still, some brief remarks are in order. Given that one is acting on a prior decision or intention that did involve the explicit or active representation of the reason, the picture of entitlement will be broadly historical and causal. It seems relevant for the sorts of cases that we're discussing that at some point in the past the explanatory consideration *was* explicitly represented in deliberation or decision-making.

¹⁰ Of course, the self-imposed standards won't be able to play their role in regulating what the agent does if they have been forgotten. But that doesn't mean that they no longer apply as standards; nor does it imply that they don't figure in the explanation of what the agent did manage to perform.

These historical conditions might then be supplemented or elaborated in terms of a cognitive architecture or infrastructure that allows for something that is not explicitly and actively represented (but at one point was so represented) to figure in the intentional description and explanation of the action. For example, I may have a reason R that is a consideration for performing some action ϕ . In response to R, I start to draw up plans for ϕ -ing, and fill them in as I proceed. But at some point, R need no longer be represented. It's done its job in deliberation, being weighed against other considerations, and leading to the decision and intention to ϕ . Nevertheless, it remains that R is the reason why I'm ϕ -ing. Intentional action is in general action that we can make (at least some minimal) sense of. And giving a reason for an action is a fundamental way of doing this, even in some cases where the reason doesn't figure in the active guidance of the performance.

4.2 THE INTERPERSONAL CASE

I think that entitlement to a reason also figures in interpersonal cases, such as that of joint action. As in the individual case, I will suggest that the structure and subject matter of reasons for certain contributions to joint or shared activity supports the claim of entitlement.

Consider the following: You and I are going to a conference downtown by car. I drive; you navigate. You say *turn left*, and I do. *Why* did I turn left? We can say:

- (6) I am turning left in order (eventually) to get downtown,

where the *in order to* locution again marks the reason for the action. But it seems to me that there is more to the structure of reasons here, that the meaning or intentionality of the action is richer than is suggested by (6). In particular, there may be a route-based consideration for turning left. For example:

- (7) I am turning left *in order to get on Dunedin in order to take High Street* to get downtown.¹¹

¹¹ Some would grant no more than that *we are turning left in order get on Dunedin in order to take High Street downtown*. But, since I'm driving, we are turning left in virtue of my turning left; arguably my turning left *is* our turning left. So even if we reject (7), we are left with a challenge very similar to what will concern me in this paper: that of understanding the intimate and seemingly explanatory and intelligibility-conferring connection between the reason involving the route-based considerations on the one hand and *my* intentional contribution to our turning left on the other.

The route-based considerations underlying your directive figure in the intentional description of my turning, *even though I do not myself have a clear sense of the route*: I may be unaware of that particular reason for which I turn. Thus, I can ask you: *now, why am I turning left here?*, and you can fill me in. Your reason sets the standard for whether I succeed at what I'm doing. And my access to the reason via communication with you is presumably what enables the reason to play the regulative role that we typically expect of the reasons *why* an agent acts.

Thus, if you mean to instruct me to *go left*, but don't speak clearly (or even misspeak) and I hear you as saying *go west*, and so turn right (which is west), that's a problem: High Street is not in that direction, and we'll have to double back or else take a different route downtown. I may be blameless, but the standard that governs and regulates what I do is given by the route-based consideration. What I was up to was getting to High Street, and I'm failing at that by turning right.¹²

These route-based considerations figure in the intentional description of the action, even though I don't have in mind the specific route we're taking. The upshot, then, is that I am *entitled to your reasons*—the route-based considerations in this case—for turning left. (Just as I'm acting for the reason that went into the deliberation that led to a prior intention.)

What would account for one person being entitled to the reasons of another? There is, presumably, a causal and historical story to be told here, as in the case of entitlement in the individual case. Moreover, it is likely that the account would advert to a cognitive infrastructure—but this time partly social. A comprehensive theory of shared agency should have something to say about this. I will not undertake that here, and will not be offering an account of this entitlement. The more modest goal of the paper is just to point to the existence of this entitlement or, at the very least, to suggest that our conception of what someone is doing in certain social contexts such as that of joint action would seem to commit us to it.

Resisting the interpersonal entitlement claim: Some may want to push back against this picture of the interpersonal case. Might we re-interpret the case so that reasons of my own make it unnecessary to invoke any entitlement to another's reasons to explain the action? For example, why isn't the reason for what I did simply that you told me to? That is:

- (8) I'm turning left in order to comply with your instructions.

¹² Compare the case of miscommunication with that of misremembering. For example, I go down the hall in order to fetch the printout, but on my way I misremember and think I'm off to check my mailbox.

I have no doubt that sometimes we do have such reasons. Perhaps I have had difficulty following instructions generally and want to work on it. Or (imagining that the navigator in this example is a woman) perhaps I have recently come to the realization that I don't take seriously instructions given to me by women, and so am making a conscious effort to follow your instructions. Or, relying on my navigational gut instincts has led me astray on many occasions and so I am now making a conscious effort to resist my intuitions and to follow someone else's lead instead. In these sorts of cases, it's reasonable to say that I'm turning left in order to comply with the instructions. But to take this as the general model threatens to distort what's going on in shared activity. Normally, following instructions is not a matter of concern and attention. The proposed re-interpretation of the case is problematic in bringing to the foreground something that is normally in the background.¹³ This is not to deny that one is *acting on the instructions*, just as one acts on a prior intention. But acting on a prior intention to φ isn't normally a *reason* for φ -ing; nor in the case of shared activity is complying with instructions to φ normally a *reason* for φ -ing.¹⁴

Let's consider another attempt at re-interpreting our case:

(9) I'm turning left in order to take advantage of your reliability as a navigator so that I can get (us) downtown.

The first thing to say here is that assuming that this is indeed a reason for me to turn left, it is quite compatible with my having other reasons as well. So if my reason for turning left is to comply with your instructions (because I think that following your instructions will get me downtown), that doesn't mean that I'm not also acting for the route-based considerations. It is easy for us to forget this because discussion of acting for a reason is often in the context of how it is that one is acting for one reason and not for another.¹⁵ One tends to forget that sometimes one might be acting for *both* reasons.

But to leave it at that is a bit misleading, because it suggests that the considerations in question are independent of one another. Recall that in our example I was suggesting that the route-based considerations that you

¹³ Admittedly, if the navigator were to ask me why I turned left, it might be natural for me to reply, "because you told me to." But the navigator presumably would only ask this if somehow there was miscommunication or confusion; perhaps she misspoke, or I misheard. I don't think that the naturalness of locutions such as (8) for what amounts to the case of a breakdown in execution should dictate what counts as the reasons for the action in the case where things proceed smoothly. I want to thank a referee for prompting these remarks.

¹⁴ Thinking of intentions as reasons also raises worries about illegitimate bootstrapping. See Bratman (1987), and Broome (2001).

¹⁵ See Davidson's central argument against Anscombe's account of reasons explanation in "Actions, Reasons, and Causes" in Davidson 1980.

have is a reason for me to turn left. The alternative suggestion for turning left is simply to make use of your reliability in this matter. But I contend that in the normal case, your reliability is not a reason for turning left. It is, rather, a consideration that rebuts a possible defeater for my entitlement to your reasons. If some question or doubt is raised about your navigational instructions, then it might be appropriate to respond by pointing out that you are reliable in giving driving directions. Thus, I might retain confidence that the reason for my turning left—whatever it is—is sound.

So it is problematic to assume that your reliability is a reason for my turning left because it distorts the normal structure of reasons in this case. It fails to show how the thought that your instructions are a good guide to going downtown can be *blocking* a potential defeater for my entitlement to your reasons. I'm not saying that your reliability cannot ever serve as a reason; it's just that in shared agency it usually is not a reason for what I'm doing. Reliability works *by* supporting the entitlement to the navigator's reasons.¹⁶

If one were to insist that the reason for my turning left is simply to comply with instructions as in (8), then we won't account for the naturalness of *looking through* the instructions to what their point is. I'm sensitive to indications of progress toward getting downtown to where the conference is. If I think we're heading in a different direction I might ask: do you realize that this is taking us north, but wasn't downtown south? Whereas if I'm just following orders, I don't need to start asking any questions, given that your instruction was simple: turn left here.

I'm not suggesting that in "looking through" the instructions I'm explicitly representing the further considerations for turning left. The whole point is that this needn't be so. But the naturalness of looking through the instructions in this sense suggests that I take the point of the instructions—the reasons underlying them—to have a significance for what *I'm* doing, for what I'm up to. They set a standard for what I take myself to be doing, and through you they can play a regulative role for how I carry out my contribution to what we're doing. I am in this way taking you as a *source* and repository of reasons for what I'm doing.

In contrast, on the proposed re-interpretation (9), I take you the navigator as, in effect, a reliable black box. Your reasons have no direct bearing on what I'm doing. I am not concerned to look through your commands to what their point is, and have no immediate interest that might be taken in the reasons underlying the directions you issue. Rather than a source of reasons for me, you (or your reliability on the matter) *constitute* my reason.

¹⁶ Burge (2013) makes something like this point in developing his 1993 account of testimony.

In this way, you (or your agency) have a status that is more like that of a reliable instrument available to me as a resource rather than a rational agent with whom I'm acting.

This is not to say that regarding the navigator in this way is necessarily wrong or problematic in general. It may be a perfectly appropriate attitude to be taken by a passenger in the back seat being driven by the main characters in our scenario. But that's because the fellow in the back is not a participant in the joint action and has no entitlement, for example, to the route-based considerations underlying the direction to turn left.¹⁷

4.3 THE LIMITS OF ENTITLEMENT

I've argued that sometimes when individuals act together, one participant may be entitled to reasons had by another. But not all reasons are available to another in this way. Some lie beyond the scope of entitlement.

The Instrumental Hypothesis. To return to the drive downtown, imagine that you the navigator have more than one reason for instructing me the driver to turn left. One is the route-based consideration. But suppose also that part of why you selected this High Street route was that it'll take us right by a café where your ex is set up outside with his laptop. You relish the opportunity to be seen with me (who wouldn't want to be seen with me?), thereby further tormenting his already tortured soul. I would never knowingly act for such a purpose (it's not nice, plus the guy is big, with a mean temper), but let us suppose that this sort of motive on your part is compatible with our acting together. (I'm imagining that you didn't orchestrate the whole ride downtown under false pretense so as to torment your ex.) So we might ask whether I am in this situation entitled to this reason of yours. Can such a consideration figure as part of the explanation for what I'm doing? That is, can we say:

(10) I'm turning left in order to drive by your ex so that we're seen by him (so that he's further tormented).

I think it's pretty clear that this consideration should *not* count as part of the reasons explanation of what I'm doing.

Part of why I am entitled to the route-based considerations as a reason for why I'm turning left is that those considerations serve an end that you and I share—namely, that of getting downtown. This suggests the *instrumental*

¹⁷ Can there be cases where I have an entitlement to your reasons, but I also have very good empirical evidence for your reliability? Having this evidence of reliability shouldn't preclude my also having an entitlement to your reasons.

hypothesis that if I am to be entitled to one of your reasons, you must regard that consideration as instrumental for or contributory to the ϕ -ing that you and I are doing together.¹⁸ Tormenting the ex in no way promotes an end we share and therefore is not a reason to which I am entitled.

If we were to adopt this simple instrumental hypothesis, it would seem that our ultimate shared end—in this case getting downtown—would define an outer boundary for entitlement. What would this say about further and *divergent* ends each of us may have for going downtown? Suppose the driver is trying to get downtown to go to the museum, and the navigator is also trying to get downtown, though not to go to the museum but rather to see a game. Would the following be a legitimate, if partial, reasons explanation?

(11) I'm turning left in order to get you to the game.

Of course, if I am aware that you are going downtown in order to get to the game, then I can easily take on as one of my ends the goal of getting you there. Then it would be uncontroversial that I turn left at least in part to further that end. But the issue is one of entitlement, and we are to imagine the possibility that I don't know why you want to go downtown. It is not so clear to me that there is entitlement to such a consideration in this case. If there is no such entitlement, then that would be some confirmation of the instrumental hypothesis regarding entitlement, since getting you to the game is not instrumental to our going downtown.

There does seem to be a difference between the instrumental route-based considerations to which I am entitled and your further end for which it's not clear that I have entitlement. However, we should not be too hasty in dismissing entitlement in this instance. A consideration in favor of entitlement stems from the thought that your further end (of getting to the game) might figure as part of the standard of success for what I'm doing, and part of what regulates how I proceed. Thus, it might constrain how I pursue the end of going downtown; how I proceed might be regulated in such a way as to go to a part of downtown that is more convenient to your end of going to the game. This might suggest that I do have an entitlement to this further end of yours, even though it is not instrumental to our explicitly shared end of getting downtown.

¹⁸ Another hypothesis would be that if X is entitled to Y's reason R as a reason for ϕ -ing, then Y must regard R as consistent with X's aims, intentions, values, etc. (This is quite different from the material about defeaters below, since it's only formulated in terms of what Y thinks.) In any case, I don't find this hypothesis to be plausible. It lets in way too much and ignores X's practical orientation.

If this is correct—and I'm not sure that it is—then the simple instrumental hypothesis about entitlement would have to be modified or supplemented to allow for some entitlement to further ends that are not instrumental to the explicitly shared end. In this way, joint activity toward some common end may, through entitlement, lead to the sharing of additional further ends—namely, those that inform how participants conceive their initial and explicitly shared end is to be brought about.

Defeaters. We've been looking at whether and how instrumental considerations affect entitlement. Let me turn now to the implication defeaters may have for entitlement. Here I will focus on an *epistemic* consideration that might be a candidate for undermining entitlement.¹⁹ Suppose you instruct me to turn left, but it's not yet clear to me that your plan is to take High Street downtown. (Maybe you haven't said why I should turn left, and maybe I'm not familiar enough with the neighborhood to glean that this is your plan.) So far, we have nothing we haven't seen before, and the suggestion has been that I'm entitled to this route-based consideration as a reason for which I'm turning left. But now suppose in addition that I know or have every reason to think that High Street is closed and that we can't get downtown this way; we can imagine that you are ignorant of this closure. How does this affect my entitlement to these route-based considerations?

If having reasons or the entitlement to reasons is like possessing epistemic justification, then we might be tempted to hold that my knowledge that High Street is blocked should defeat or undermine entitlement to the route-based consideration for turning left. But it's not clear that entitlement to reasons for action works quite like this. That's because entitlement has been understood here as concerned with the reason for which one acts, and acting for reasons seems to be quite different from possessing epistemic justification. For example, I may know that there is no beer in the fridge, so I have no justification for thinking that I can get beer from the fridge. Nevertheless, momentarily forgetting that there is no beer in the fridge, I open the fridge *in order to get a beer*.²⁰ In the individual case, then, possessing a defeater for a reason is compatible (at least on some occasions) with acting for this reason. Likewise, in the collective case, it seems possible for me to act for your reason, even though I possess defeaters for it. I possess a defeater for thinking that High Street gives us access to downtown. Thus, the reason for turning left is

¹⁹ Non-epistemic considerations might similarly undermine entitlement. This is implicit in the earlier case of driving by the café to torment the ex.

²⁰ I'm working here with the notion of a mental state defeater that is available to the agent upon reflection. The agent does possess the defeater understood in this way. Since the agent temporarily forgets, they do not possess a defeater understood as a *conscious* mental state defeater.

undermined; it's no longer a reason to turn left. Nevertheless, it appears that I may still be turning left for that reason, since I don't yet realize that the point of turning left was to get on High Street. Possessing a defeater therefore doesn't necessarily block entitlement.

But sometimes the possession of a defeater for a reason is such that it is hard to make sense of acting for that reason. In the beer example, I temporarily forget that the fridge is empty of beer, and in that case it is easy enough to imagine that I open the fridge in order to get a beer. But suppose that I don't suffer anything like a memory lapse, undue distraction, or some other *glitch*—by which I mean a cause of error that doesn't put in question my rational agency or the intelligibility of acting for that reason.²¹ Imagine instead a case that is non-glitchy, but more deeply weird: I remember that the fridge is empty—that is, I actively recall this fact—but then go on to open the fridge *in order to get a beer*. I don't think that the action can be made sense of in this way. I might, of course, *pretend* to open the fridge to get a beer. Or I might open the fridge in order to get something else. Or, I might open the fridge in order to check to see whether my memory is playing tricks on me. But I can't just open the fridge in order to get a beer. So, sometimes the presence of the defeater for a reason is such that we cannot make sense of the agent being able to act on this reason. In light of this, I suggest that, *absent glitches*, possessing a defeater for reason R for one's φ -ing would undermine one's entitlement to R. That is, there is no entitlement to R in circumstances where one has a defeater for R, and φ -ing for R would compromise one's rational agency.

Return to the case where I know or have good reason to believe that downtown is inaccessible via High Street. Whether I have an entitlement to the route-based considerations in this instance (i.e. whether I turn left in order to take High Street downtown) will depend on why the defeater I possessed did not alter the course of action. Was it due to some mere glitch in communication or memory? Or does it reflect a more serious compromising of my rationality or rational agency? Suppose I simply forget about the blockage on High Street. Or maybe I'm unlucky and so unfamiliar with the area that it doesn't occur to me that you are directing us to High Street. Or maybe there is miscommunication and I am under the impression that you are having us take a different route. In all these cases my failure to bring to

²¹ The idea is that suffering from this sort of error doesn't put in question one's rationality. Of course, some serious memory deficits do compromise one's rationality and can in some circumstances put in doubt the subject's ability to act for a certain reasons. The idea of a glitch is related to Burge's notion of brute error. But, as Burge (1988: 657) characterizes it, "brute errors do not result from any sort of carelessness, malfunction, or irrationality" on the part of the subject. Although glitches are compatible with rationality, arguably they do involve some form of malfunction.

bear the defeater should, I think, count as a glitch; it doesn't reflect a deeper problem with my rational agency. It's plausible, then, that in following your instruction to turn left, I am doing so in order to take High Street. So here we seem to have entitlement despite the presence of a defeater.

Suppose instead that there are no glitches. Although you don't give the reason for turning left, there is no miscommunication that suggests that we're not taking High Street. I remember that High Street is blocked. Moreover I am familiar with the neighborhood, and am not distracted from what we're doing, and so it would not be hard for me to recognize that the plan is to take High Street. And yet, ignoring all the signs, I follow your instructions even though I know that High Street is blocked. This, I submit, puts in question the intelligibility of turning left *in order to* take High Street downtown. Here, the defeater *does* seem to undermine entitlement.²² I conclude that the possession of defeaters, suitably understood, offers another important constraint on the interpersonal entitlement to reasons.

CONCLUSION

We started with the thesis that sometimes the reasons that explain one's action and account for its status as intentional are not one's own (in the sense that one does not represent those reasons for oneself). More specifically, one's entitlement to the reasons of another enables those reasons to figure as the reasons for which one acts. The claim will no doubt strike many as controversial. For example, it challenges some conceptions of the agent's non-observational epistemic access to the reasons for which she acts. In these cases of entitlement, this access would likely instead have to be understood more along the lines of testimony and whatever sort of warrant for belief we acquire through testimony. There are also implications for discussion of responsibility. A condition for responsibility and blame for an action is its being attributable in some robust sense to the agent. On some views, this is

²² Another case is where the navigator misleads me on purpose. I may have reasons for thinking that taking High Street should be ruled out or is much less desirable than some alternative route. If I realized that the plan was to take High Street, I would have insisted on some alternative. But you manage to trick or otherwise manipulate me so that I fail to glean what the plan is. I am tempted to think that what happens to me in this case is not the equivalent to a glitch but amounts, rather, to a deeper undermining of my rationality. It therefore would preclude entitlement. If one is not tempted to think that this form of trickery amounts to an undermining of rationality, then we might point instead to a further restriction on entitlement. We might say that absent glitches *and deception*, possessing a defeater for some reason R for one's ϕ -ing would undermine one's entitlement to R.

understood in terms of agent's ability to give some accounting for what he has done in terms of his evaluative judgments.²³ But it appears that in cases of entitlement, action is attributable to the agent even though he would only be able fully to account for his φ -ing with help from someone else (such as the navigator in the example of the drive downtown). I have no doubt, then, that the central claim of the paper is provocative and work is needed on a number of issues to accommodate it. But it becomes easier to think that one individual can be entitled to the reasons of another once we recognize the scope and limits of entitlement, which I've only begun to sketch in the latter part of the paper. Properly understood, entitlement becomes a plausible and intriguing way to think of how individuals in certain social contexts such as that of shared activity can be related to one another—namely, in terms of the reasons of one explaining the actions of another.²⁴

References

- Anscombe, G. E. M. (1957). "Intention." *Proceedings of the Aristotelian Society* New Series 57: 321–32.
- Anscombe, G. E. M. (1963). *Intention*, 2nd edn. (Ithaca: Cornell University Press).
- Bratman, Michael. (1987). *Intentions, Plans, and Practical Reason*. (Cambridge, MA: Harvard University Press).
- Broome, John. (2001). "Are Intentions Reasons? And How Should We Cope with Incommensurable Values?" in *Practical Rationality and Preference: Essays for David Gauthier*, edited by C. Morris and A. Ripstein, 98–120. (Cambridge: Cambridge University Press).
- Burge, Tyler. (1988). "Individualism and Self-Knowledge." *Journal of Philosophy* 85: 649–63.
- Burge, Tyler. (1993). "Content Preservation." *The Philosophical Review* 102: 457–88.

²³ Scanlon 2008, Smith 2005. See Shoemaker 2011 for critical discussion.

²⁴ Earlier versions of this paper were presented at the New Orleans Workshop on Agency and Responsibility (NOWAR) at Tulane University, and at the Collective Intentionality IX conference at Indiana University as part of a symposium on social explanation. I have benefitted from audience comments on those occasions, as well as from seminar participants at Ohio State University. Thanks also to Michael Bratman, Kirk Ludwig, Sarah Stroud, Facundo Alonso, Vida Yao, Raimo Tuomela, Lisa Downing, and William Taschek. This paper was also presented more recently at the Workshop on Joint Practical Knowledge in Santiago, Chile (organized by Johannes Roessler (Warwick) and Glenda Satne (Alberto Hurtado University)), and at a symposium session at the Eastern APA (with comments by Kirk Ludwig and Max Lewis). I received very helpful feedback on those occasions, but it was too late to take them up here. I hope to address these issues on a future occasion.

- Burge, Tyler. (2013). "Postscript to Content Preservation." In *Cognition Through Understanding*, 254–84. (Oxford: Oxford University Press).
- Davidson, Donald. (1963). "Actions, Reasons, and Causes." *Journal of Philosophy* 60: 685–700.
- Davidson, Donald. (1980). *Essays on Actions and Events*. (Oxford: Oxford University Press).
- Frankfurt, Harry. (1978). "The Problem of Action." *American Philosophical Quarterly* 15: 157–62.
- Harman, Gilbert. (1986). *Change in View*. (Cambridge, MA: MIT Press).
- Jeannerod, Marc. (2006). *Motor Cognition: What Actions Tell the Self*. (Oxford: Oxford University Press).
- Kavka, Gregory S. (1983). "The Toxin Puzzle." *Analysis* 43: 1, 33–6.
- Scanlon, Thomas. (2008). *Moral Dimensions: Permissibility, Meaning, Blame* (Cambridge, MA: Harvard University Press).
- Shoemaker, David. (2011). "Attributability, Answerability, and Accountability: Toward A Wider Theory of Moral Responsibility." *Ethics* 121: 602–32.
- Smith, Angela. (2005). "Responsibility for Attitudes: Activity and Passivity in Mental Life." *Ethics* 115: 236–71.
- Wilson, George. (1989). *The Intentionality of Human Action*. (Stanford, CA: Stanford University Press).