LUND UNIVERSITY

Practical Knowledge and Acting Together

Blomberg, Olle

Published in:
Socially Extended Epistemology

DOI:
10.1093/oso/9780198801764.003.0006

2018

Document Version:
Early version, also known as pre-print

Link to publication

Citation for published version (APA):
Blomberg, O. (2018). Practical Knowledge and Acting Together. In J. A. Carter, A. Clark, J. Kallestrup, O. Palermos, & D. Pritchard (Eds.), *Socially Extended Epistemology* (pp. 87-111). Oxford University Press. https://doi.org/10.1093/oso/9780198801764.003.0006

Total number of authors:
1

General rights
Unless other specific re-use rights are stated the following general rights apply:
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.
• Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
• You may not further distribute the material or use it for any profit-making activity or commercial gain
• You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: https://creativecommons.org/licenses/

Take down policy
If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Penultimate draft.

Forthcoming in *Socially Extended Epistemology*, edited by J. A. Carter, A. Clark, J. Kallestrup, O. Palermos, and D. Pritchard, OUP 2018.

Please quote from the published version.

Practical knowledge and acting together

Olle Blomberg

olle.blomberg@gmail.com

According to one influential philosophical view of human agency, for an agent to perform an action intentionally is essentially for her to manifest a kind of self-knowledge: An agent is intentionally φ-ing if and only if she has a special kind of practical and non-observational knowledge that this is what she is doing. I here argue that this self-knowledge view faces serious problems when extended to account for intentional actions performed by several agents together as a result of a joint decision. This suggests that practical and non-observational knowledge is not essential to intentional action as such, since a theory of intentional action ought to be able to make sense of such paradigm cases of joint intentional action as well as cases of singular intentional action. Existing attempts to extend the self-knowledge view to joint intentional action face an unfortunate trilemma. These attempts must (i) require that participants have non-observational knowledge of each other's non-observational knowledge, which in turn requires that the participants are either really one and the same agent, or that they are, for all relevant purposes, clones of each other; (ii) limit their explanatory scope to intentional actions performed by groups with a single decision maker at the top of an organisational hierarchy, which arguably are not really *shared* or *jointly* performed; or (iii) require that the intention of the group is the result of a public representational act that the group members will have observational knowledge of in a way that is incompatible with a core commitment of the self-knowledge view. I argue that each of these options is unacceptable given a self-knowledge view of intentional action.

Word count: 14,265 words

1. Introduction

According to an influential view in contemporary philosophy of action, intentional action is essentially a manifestation of an agent's practical and non-observational knowledge of what she is doing intentionally (see, e.g., Anscombe 1969; Velleman 1989; Setiya 2008; Rödl 2007; Thompson 2008). On this *self-knowledge view* of intentional action, an intention to φ embodies a cognitive commitment that one is φ-ing or that one will φ. One consideration in favour of this view is that we typically express an intention to φ in the same way that we express a belief-like commitment about what we are doing or will do: "I'm φ-ing" or "I'm going to φ". But this commitment is practical rather than theoretical in that it causes what it represents as true rather than being caused by it. Another consideration that might be seen as favouring the view is that, if intention didn't embody this cognitive commitment,

then we would be at a loss to explain why it is irrational to form conflicting intentions or in other ways flout principles of means-end rationality (Velleman 2007; but see Bratman 2009). After all, if I normally did not believe (or had some other belief-like cognitive commitment) that I would be doing what I intended, then it appears there would be no fault in intending to do what I know is not co-realisable with something else that I already intend to do. But it is a rational mistake to have such conflicting intentions, so we must presumably believe that we are doing or will do what we intend. Furthermore, these cognitive commitments must arguably provide me with knowledge of what I am doing intentionally when I am acting (perhaps in virtue of my skills and know-how, see e.g. Setiya 2008; 2009; O'Brien 2003; Grünbaum 2009; Small 2012). That the commitment is true is not sufficient for intentional action. My chopping of an onion will not be intentional if the cognitive commitment embodied in my intention to chop it is only accidentally true. An intentional action must arguably be a "manifestation of [the agent's] competence" (Sosa 2015, 23, n. 24).

In contrast to an agent's knowledge of what she is doing unintentionally or of what other agents are doing, an agent's knowledge of what she is doing intentionally is non-observational. If you ask me why I am chopping an onion and I'm not aware of the fact that I'm chopping it or only become aware of it by observing myself, then I would not be the source and creator of this action. The chopping of the onion would then be something that happened to me. If it is an action that is performed intentionally, then I can answer that I'm chopping an onion because I'm making dinner, and do so without recourse to observation. This is not to deny that successfully performing an action φ often depends on the availability of perceptual information that the agent uses to guide and monitor her action performance, but the agent doesn't rely on this information to generate and test hypotheses about what it is that she is intentionally doing (see Falvey 2000). For example, I need to track where on the cutting board the onion is and how it is spatially located with respect to my hand and the blade of the knife. The perceptual information does not provide the justification for a belief about what it is that I am doing.

The self-knowledge view is controversial when it comes to intentional actions that are performed by a single agent, but I think proponents of the view have successfully shown that various alleged counterexamples fail to decisively exemplify intentional action without practical or non-observational knowledge (see e.g., Setiya 2008; Velleman 1989 ch. 4; Small 2012, sect. 5; Stathopoulos 2016). I will not discuss such purported counterexamples. Instead, I will argue that the self-knowledge view runs into deep problems when faced with accounting for intentional actions that are performed by several agents together. A theory of intentional action should be able to throw light on such joint intentional actions in addition to the singular intentional actions that are more familiar in philosophy of action.

Consider the intentional action of making dinner, but this time performed by you and me together. In this case, some component actions of this larger action are performed by me, others by you. There is, I think, no reason to take the singular

case, where all the component actions are performed by one agent, to be more fundamental than this joint case. Many of the complex actions that we perform on a daily basis are performed together with others, and many of the actions we perform on our own are typically learnt at some point and first performed as joint intentional actions with co-participants or teachers. I didn't begin brushing my teeth on my own. At one point it was a collaborative endeavour involving me and a parent. In light of this, privileging singular intentional action is not an innocent and convenient starting point, but may well introduce distortions and biases in one's theorising about human agency. This is not to say that single human agents intentionally moving their bodies is not in some sense a fundamental starting point of human agency. But once such simple actions are combined into larger complex wholes such as making dinner or brushing one's teeth, there is little reason to think that the singular cases are more fundamental than the joint cases in a way that matters for the philosophy of action.

In the last twenty years, a few attempts have been made to extend the self-knowledge view to account for joint intentional actions (Velleman 1997; Stoutland 2008; Laurence 2011; Rödl 2015; Schmid 2016).[1] According to these extensions, we would arguably jointly intentionally φ if and only if we had "joint practical knowledge" or "joint non-observational knowledge" that we were φ-ing. (From now on, I will interchangeably use 'practical knowledge' and 'non-observational knowledge' to refer to the knowledge that an agent has of her intentional actions according to the self-knowledge view. Others have also used the terms 'agent's knowledge' and 'agential knowledge' to refer to this type of knowledge.) This is an explicit commitment for Stoutland (2008), Laurence (2011) and Schmid (2016).[2] While neither Velleman nor Rödl explicitly discusses joint practical knowledge, both are implicitly committed to the possibility of such knowledge given their views of intention and intentional action (see Velleman 1989; Rödl 2007). The extensions are supposed to apply to small-scale egalitarian cases of joint intentional action and the actions of groups with authority structures and the potential for changing membership. Examples of the former include two people going for a walk together (Velleman 1997); two people playing chess or nailing a long board together (Stoutland 2008); a band of robbers carrying out a heist and several men pushing a car to a gas station together (Laurence 2011); two people painting a house together (Rödl 2015); and two people writing a paper together (Schmid 2016). Examples of the latter include a navigational manoeuvre made at sea by a ship's crew (Laurence 2011), a Philosophy Department filling a vacancy in its ranks by hiring someone (Velleman 1997), and the actions of corporations (Stoutland 2008).

---

[1] Seemann (2009) glosses his account of joint intentional action as "Anscombian", but it is not an extension of the self-knowledge view.

[2] However, Schmid's main concern is a pre-reflective awareness of *who* the subject of the action is ("we", not "I") rather than knowledge of *what* the group is doing.

In this chapter, I will critically examine these attempts to extend the self-knowledge view and articulate some general challenges that any extension of the view has to face. The plan is as follows: In section 2, I examine why an extension of the self-knowledge view cannot build on a purely distributive conception of joint practical knowledge. On such a conception, joint practical knowledge is nothing more than the practical knowledge of each participant that they are jointly intentionally φ-ing. This would require participants to have non-observational knowledge of each other's non-observational knowledge of what they are jointly intentionally doing, which would appear to be impossible unless the participants were one and the same agent or, for all practical purposes, were clones of each other.

In section 3, I examine whether and how a single token intention held by several agents together as a group could provide them with joint practical knowledge. I do this by drawing on two accounts of how individual intentions provide individual agents with practical knowledge: a reliabilist interpretation of Anscombe's (1969) account and Velleman's (1989; 2015) evidentialist account. I argue that, while there is nothing especially problematic in making sense of how a belief-like commitment embodied in a group's intention could be justified, it is hard to make sense of the claim that the group members could have non-observational knowledge of a single jointly held token intention. In particular, I argue that the most detailed and carefully worked out extension of the self-knowledge view to joint intentional action (due to Velleman 1997) fails to make sense of this.

In section 4, I consider whether a more robustly collectivist conception of joint practical knowledge might be more successful. On such a collective conception, joint practical knowledge is the practical knowledge of the group that they, the group members together, are jointly intentionally φ-ing. Here, I show that attempts by Stoutland and Laurence to explicitly illustrate the possibility of such collective joint practical knowledge are unsatisfactory. Stoutland and Laurence draw on examples where there is a single individual agent that is calling the shots at the top of an organisational hierarchy. They take these to be examples of a collective agent performing an intentional action, but given the self-knowledge view, they are rather examples of a single individual agent performing an action. An individual agent's action possibilities can be radically extended by chains of command and relations of authority that enable the agent to make decisions on behalf of a whole group or organisation. These are interesting cases of social actions, but given the self-knowledge view, they do not involve several agents together performing a *shared* or *joint* intentional action. Nor would they be cases of shared or joint intentional action if the organisation as a whole were a distinct single agent with non-observational knowledge of what *it* was doing intentionally. This would be an exercise of ordinary individual agency, but realised at the group level. What the proponent of the self-knowledge view needs to do is to explain how several group members can each have non-observational knowledge of what they are jointly intentionally doing. But this brings us back to the problems with the distributive conception of joint practical knowledge that were presented in section 2. I conclude by considering what general

lessons we can draw from the problems faced by an extension of the self-knowledge view to joint intentional action.

2. Distributive joint practical knowledge

Since joint intentional action undeniably involves several participants, a natural first thought might be that joint practical knowledge is simply the sum of each participant's practical knowledge of their common single joint action. On this distributive conception, if we are jointly intentionally going for a walk together, then each of us must know without observation that we are walking. This knowledge is practical; it causes our walking rather than being caused by it. Let us see what can be made of such a suggestion.

According to several accounts of practical knowledge, an agent only has practical knowledge of the basic action that she is currently performing, but she does not necessarily know that she is bringing about the more distal goal that she is aiming at (e.g. O'Brien 2003; Setiya 2008). While an archer may not have practical knowledge that she is shooting an arrow that will hit the bull's eye, she will normally have practical knowledge that she is performing the basic action of relaxing the fingers on her string hand and releasing the arrow. An action is a basic action for an agent if and only if the agent can intend to perform it without intending to perform any other action by means of which she intends to perform it (Searle 1980, 65–66). It is an action that the agent can perform "just like that" without the need for any further practical reasoning. Given that a basic intention—that is, an intention to perform a basic action—embodies a cognitive commitment that the agent is performing the intended action, then it is plausible that the agent will non-observationally know in virtue of her skill and expertise that she is performing it.

I trust that it is at least intuitive that the commitment in a basic intention would be justified. According to some interpreters, Anscombe has a reliabilist account of this justification, according to which a reliable connection between intending to φ and actually φ-ing is sufficient for practical knowledge (see Velleman 2007; Grünbaum 2009).[3] Given this reliable connection, the agent's non-observational knowledge of her own intention can also provide her with non-observational knowledge of what she is doing intentionally. If the intended outcome isn't brought about, then there is of course no non-observational knowledge. But in that case, neither is the outcome intentionally brought about.

However, on this reliabilist picture, practical knowledge would be divorced from the agent's own point of view in a way that is in tension with the intuitions that make the self-knowledge view attractive in the first place. A third-person spectator who knew the content of the agent's intention and who knew that the intention's

[3] For a related view of intentional action, see (Sosa 2015, ch. 1).

connection to the intended action was reliable would be in almost the same epistemic position as the agent herself with respect to her intentional action. Practical knowledge is supposed to be a form of "insider" knowledge, which is distinct from how we know what others are doing or from how we know what we ourselves are doing unintentionally. Partly to avoid this problem, Setiya (2008; 2009) argues that the agent is entitled to the cognitive commitment embodied in an intention by virtue of *knowing how* to perform the basic action(s) she is performing or will perform. This is supposed to go beyond merely having a reliable disposition to perform the action(s), and it is thereby supposed to explain how practical knowledge can be a form of insider knowledge.

Whatever justification the self-knowledge view appeals to, a basic action-based account of practical knowledge may seem to be readily extendable to cases of joint intentional action. After all, if an action that an agent can perform "just like that" with another agent were learned and practiced together with others, then it is plausible that the basic intention to perform that action could specify not merely the agent's own bodily movement but also the other agent's bodily movement in its content (Blomberg 2011). There can be reliable connections between basic intentions and intended outcomes that are jointly brought about, and the knowledge of how to perform the action may only become manifest when co-exercised with another agent's knowledge thereof. Consider shaking hands with someone or performing a joint move in a dance with a partner with whom you have practiced extensively. In such a case, you and the other can arguably each perform the action spontaneously, without the involvement of any practical reasoning. If you were to try to faithfully perform your own part of the joint movement in the other's absence, then you would probably fail.

A basic joint intentional action, as we might call it, would be a joint φ-ing that is caused and coordinated by each participant's basic intention to φ. Each participant would non-observationally know that they were φ-ing in the sense that there would be a joint bodily movement φ that each was performing. What we get are two agents who each have practical knowledge that is socially extended in the following sense: The knowledge concerns not only the agent's own bodily movement but also that of the others. The states of knowledge of the participants would have a common object, namely one and the same joint bodily movement. In such a case, one might perhaps say that the participants would have joint practical knowledge that they were φ-ing.[4]

However, it would be misleading to call this *joint* practical knowledge. In the individual case, the self-knowledge view requires not only that the agent non-

---

[4] In analogy with Setiya's (2008) account of singular practical knowledge, each participant could, perhaps, be entitled to the joint practical knowledge in virtue of his or her joint know-how or in virtue of all the participants' "group know-how" (see Palermos and Tollefsen, this volume).

observationally knows that she is φ-ing. To intentionally φ, the agent must also non-observationally know that she is φ-ing intentionally.[5] If the agent didn't also know this, then non-observational knowledge would be just another way of knowing of something that, as far as the agent non-observationally knows, might have been done unintentionally. Hence, an extension of the self-knowledge view should require that each has non-observational knowledge that he or she is φ-ing jointly intentionally with the others. All versions of the self-knowledge view that I am discussing in this chapter are committed to the claim that participants must have non-observational knowledge not only of what they are jointly doing, but also of what their shared intention is. Given a distributive conception of joint practical knowledge, this means that each participant must have non-observational knowledge of the others' non-observational knowledge. But to know that the φ-ing with the others is *jointly intentional* or, to put it differently, to acquire knowledge of what the others non-observationally know, each has to rely on observation and inference to acquire knowledge of the others' basic intentions (Bratman 2014, 58). In fact, each would also need observational knowledge of the others' observational knowledge of her own non-observational knowledge, observational knowledge of the others' observational knowledge of her own observational knowledge of the others' non-observational knowledge, and so on and so forth. In short, the agents would need to have common knowledge of their intentions to φ with the other in order to know that their φ-ing was jointly intentional (see Blomberg 2016, sect. 5). But this would not be joint non-observational knowledge. Given the self-knowledge view, we would at best end up with a case of two overlapping socially extended actions that unfold under conditions of common knowledge. This would not qualify as a case of joint intentional action.

Furthermore, this basic action-based account of joint practical knowledge rests on a misleading analogy between singular intentional action and joint intentional action. For an agent to perform a singular intentional action, there must be a basic action by which she performs that action at any one time (unless this action is itself a basic action). And for several agents to perform a joint intentional action, each agent must perform some basic action (in some cases, this may be a single basic joint action of which each is an agent). But for several agents to perform a joint intentional action, there need not be any basic *joint* action that they perform.[6] No basic joint action

---

[5] On some theories of intention, this will be a consequence of the fact that the intention that the agent is acting on has as part of its content that the intention itself causes the represented behaviour (see e.g. Searle 1980; Velleman 1989).

[6] At least not in the sense of 'basic joint action' under consideration. Herbert Clark uses the term 'basic joint action' differently to refer to the most general conception of a joint action that is shared among participants, such as "assembling the TV stand" for example (2006, 129). In this usage, the term refers to something like the shared intention *with which* the participants are acting and coordinating (Anscombe

needs to be performed in examples such as two people walking together, writing a paper together, or painting a house together—cases that are paradigm cases of joint intentional action (although other paradigm cases, such as two people dancing the tango, arguably do involve some basic joint actions). In this respect, joint intentional actions are more akin to complex singular intentional actions. In a complex singular action such as that of my making dinner on my own, there is no temporally extended basic action by which I am making the dinner. Rather, the larger action of making dinner is composed of many smaller component intentional actions that are performed during different phases. Hence, it is not a basic action that provides the unity of the larger action, but rather the agent's intentions and practical reasoning throughout the duration of the action. Similarly, what provides the unity of the larger joint intentional action is the agents' intentions and practical deliberations.

This more appropriate analogy between joint and singular action suggests another way of making sense of a distributed conception of joint practical knowledge.[7] This strategy would emphasise the so-called "openness" of the progressive verb form, that is, emphasise that it is not the case that "*A* is [or was] φ-ing" implies that "*A* [will have] φ-d" (for accounts of practical knowledge grounded in the failure of this inference, see Falvey 2000; Thompson 2008; Small 2012). When I am chopping an onion as a component of the larger action of making dinner, the larger action is underway as I am performing the basic action of chopping the onion. But the larger action is not yet completed and, arguably, may never be. After all, there might be a prolonged power outage in my neighbourhood. Or I might stumble in the kitchen and end up hospitalised before I am finished. Even if such an accident occurs—and even if all of them *could* occur—this doesn't undermine the fact that I was nevertheless making dinner prior to the accident's occurrence. Hence, it need not undermine my knowledge that I was making dinner. The making of the dinner was earlier underway, and I wasn't wrong about that. I can non-observationally know that I'm making dinner as I am chopping the onion in spite of various possible but unlikely accidents that would stop my progress. One might think that I can also know this in spite of the possibility, unlikely or not, that I will change my mind, abort the cooking before completion and instead have dinner at a nearby restaurant.

However, if my action is not completed because of a failure rather than my change of mind, and this failure is not accidental but due to my lack of know-how or necessary tools, or due to the presence of another agent who is intentionally trying to stop me, then the cognitive commitment embodied in the intention may indeed be unjustified. The failure would not be accidental at all and, in such circumstances, a

---

1969, sec. 26). But this is not the sort of basicness that I have in mind when using the term.

[7] Indeed, joint intentional action and complex singular intentional action are arguably not merely analogous in this way. They are two flavours of one and the same phenomenon of complex intentional action (see Rovane 1998, ch. 4).

lucky successful match between intended outcome and actual result would be merely co-incidental. In such a case, my intention would not provide me with knowledge—and a lucky outcome would arguably not be intentional (see Small 2012, sect. 4; Sosa 2015, ch. 1).

Now, what I do intentionally will often depend on what other agents are doing intentionally in a way that is consistent with my having non-observational knowledge of what I am doing. For example, my non-observational knowledge that I'm going to Golden Gate Park depends on the taxi driver's intentional action of driving the car we are in to Golden Gate Park. As in the case where I am chopping onions, whether I have non-observational knowledge of what I am doing in this case is not a matter of whether the action will be completed successfully (the taxi might crash and go up in flames), but rather a matter of whether I non-observationally know that I'm involved in a process that is in progress. If I can have non-observational knowledge that I'm making dinner when I'm chopping the onions, then I can also have non-observational knowledge that I'm on my way to Golden Gate Park as I'm stepping into the taxi. My dependence on another agent in this case is not especially problematic for an account of practical knowledge. Can the proponent of the self-knowledge view appeal to these considerations to make sense of joint practical knowledge?

Unfortunately, this is not the case. Whether I'm engaged in a joint intentional action with you does not only depend on a generally cooperative social order that we both inhabit. It also depends on the conclusion of your practical reasoning. Furthermore, we are arguably each aware that this is so. While I can non-observationally know that I'm making dinner even if a power outage will stop my action from reaching completion, I cannot non-observationally know that I'm making dinner with you if you never decide to join me! To know that I'm making dinner with you, I need to be informed of your choice to participate, and this knowledge will be partly justified by my observations and inferences ("partly", because it will also depend on my non-observational knowledge of my own intentions and actions). For example, I might predict that you will choose to make dinner with me if you are informed of my intention to make dinner with you (see Bratman 1997). This is different from my dependence on the taxi driver, who I only depend on to perform a certain expected role in a social system (Schuetz 1953). I take the taxi driver to be a part of my taken-for-granted social environment, rather than as a co-participant in a joint intentional action. He or she is related to my going to Golden Gate Park in somewhat the same way as the knife, the cutting board and the onions are related to my making dinner. I need not observe either of them to justify my judgement about what I am doing intentionally (making dinner, going to Golden Gate Park), even if I need to take note of them in guiding and monitoring the progress of my action.

To conclude, for each participant in a joint intentional action to know what *they* are doing jointly intentionally, each must rely on observationally and inferentially based knowledge about the intentions of the others. For each participant to non-observationally know that they are φ-ing jointly intentionally, part of what each would have to know non-observationally would be that the other non-

observationally knows that they are φ-ing jointly intentionally.[8] But such interpersonal reflexivity is impossible unless the participants are all either (a) really one and the same agent or (b) identical clones with common knowledge of the peculiar fact that they have (for all relevant purposes) the same preferences, thoughts, and reasoning processes. We can set aside the first possibility since it would then no longer be a case of joint intentional action, but one might think that there are realistic cases that come close to (b). Suppose that you and I are buying a house together. We know each other extremely well, and we have extensively discussed what kind of house we want to buy. At a meeting with the broker, we are faced with an offer that perfectly matches our shared preferences. In this situation, one of us could simply go ahead and say: "We are buying it!". Suppose that each of us knows that his or her own decision will reflect the other's and this is common knowledge between us.[9] Wouldn't the speaker here express the practical knowledge we both have? I think not. It would be a misleading way to describe what would be going on. Either we have already made a conditional decision together to buy a house that fits certain criteria—a decision that each has knowledge of through observation—and one of us is merely acting on this prior tacit or explicit decision, or else the statement is not expressing a decision that is *ours* but rather the speaker's own unilateral decision. In the latter case, the decision is expressed with the use of an aspirational "we", that is, with the encouragement and hope that the other will not protest and dissent from the decision (see List and Pettit 2011, 194). Arguably, articulating a plausible account of joint practical knowledge on the distributed conception is bound to fail.

---

[8] The phrase "they are φ-ing jointly intentionally" is somewhat cumbersome, but it is preferable to "they are jointly φ-ing intentionally". This is because the latter might be understood as something that is intentionally done by only one or several agents, but which is nevertheless not done jointly intentionally. For example, suppose that you are standing still on a trampoline and I jump onto it with all my weight with the intention that we bounce off the trampoline; then it seems right that, as far as I am concerned, we are jointly bouncing off the trampoline intentionally. Furthermore, suppose that you predict that I'm going to jump onto it, and you prepare for the bounce by bending your knees slightly to facilitate the joint bouncing, also with the intention that we bounce off the trampoline. In this case, as far as each of us is concerned, we are jointly bouncing off the trampoline intentionally. But it nevertheless seems wrong to say, in light of the self-knowledge view, that we are bouncing off the trampoline jointly intentionally, since our intentions are not properly connected or unified.

[9] Thanks to Nathan Hauthaler for this example.

3. Joint decisions and practical knowledge

The root of the problems with the distributive conception of joint practical knowledge is that, on this conception, each participant has his or her own intention. In effect, each must also have his or her own private joint practical knowledge of what they are doing together. Perhaps in order to avoid the problems associated with the distributive conception, Velleman, Laurence and Rödl all hold that there is a single token intention that the participants of a joint intentional action have together, as a group.[10] In Velleman's account, this single token intention is a public representation that in some sense belongs to all the participants, in Laurence's account it is "a single rational order" and, in Rödl's sketch of an account, it is a single action-governing "principle". Schmid seems to hold a similar view since he writes that the members of a social agent "are co-authors of their attitudes in quite a literal sense" (2016, 71–72). Now, one way in which a single agent forms an intention and gains non-observational knowledge of what she is doing is by making and acting on a decision. Perhaps we can make sense of several agents together having joint practical knowledge in virtue of making a joint decision and thereby forming a single intention that belongs to all.

When a group of people makes a decision, there will be some public representation or public representational act by virtue of which the decision is made. Perhaps a group can also simply acquire an intention without its being formed through a joint decision in which no such public representational act is involved. For example, a collection of people on a beach whose attention is captured by a swimmer's cry for help might acquire an intention to rescue the distressed, without any kind of agreement or decision being made. But cases in which a group acts on an intention that is formed through a group decision represent one kind of paradigm case of joint intentional action. This kind ought to be captured by an extension of the self-knowledge view. At any rate, it is this sort of case that I will focus on in this chapter.

When we apply accounts of practical knowledge to shared intentions, two main issues arise. First, are there any novel problems (or solutions) related to the justification of the cognitive commitment that is embodied in the group's intention? I address this in subsection 3.1. Secondly, and more generally, how can the group members together have non-observational knowledge of their intention? I address

---

[10] There is a sense in which accounts that reduce shared intention to a distributed pattern of individual intentions also take there to be a single token intention that is held by all the parties together, as a group (see Bratman 2014, 13, 131). But in such a reductive account, each party can nevertheless instantiate his or her own part of the pattern even in the absence of the parts of the others. In contrast, I take it that this is not a possibility on Velleman's, Laurence's and Rödl's views, according to which the single shared intention cannot be reduced to a pattern of ordinary individual intentions which are distributed among the parties.

this in subsection 3.2 through a detailed discussion of Velleman's extension of the self-knowledge view.

3.1 The justification of the joint cognitive commitment

For illustrative purposes, let us consider a toy example of an executive board of a corporation that makes a decision that the employees then execute (the example can be considered as an elaboration of an example provided by Stoutland (2008, 546) to illustrate the possibility of joint practical knowledge, see section 4). Suppose that the corporation is in the home furnishing retail business and the board is considering whether to open five stores in a country where they currently do not have a presence. After much discussion concerning profits and costs, risks associated with involvement in the country's corrupt local politics, and the compatibility of various courses of action with the corporation's expressed aims and core values, the board decides to make contact with the country's authorities in order to go ahead and open the stores. We can imagine various procedures by which the decision is made: Perhaps the discussion leads to a consensus among the board members regarding what would be best for the corporation to do. Or perhaps the decision is made by majority voting. Regardless of the decision procedure used, the decision is established when the chairman of the board utters "We hereby decide to open the stores". As a result of this, the corporation successfully opens the five new stores.

The question is how the cognitive commitment embodied in the intention formed by the corporation can be justified so that the corporation or the board has non-observational knowledge that it is opening the five stores (or at least that it is contacting the country's relevant authorities). The reliabilist account of practical knowledge that I briefly presented in section 2 can be relatively unproblematically scaled up to the level of this kind of corporate agent. After all, there is nothing puzzling about there being reliable connections between executive decisions and intended results in an efficient and robust organisation. Furthermore, the reliabilist account's failure to make sense of the suggestion that practical knowledge is a form of insider knowledge looks like an advantage when the account is extended to the agency of groups or organisations. After all, there will in many cases be no special insider knowledge of the processes that connect a group's decision or intention and the ensuing action of the group or organisation. Suppose that another corporation with vested interests in the home furnishing retail business has planted a camera and a microphone inside our corporation's boardroom. This unscrupulous competitor also knows a lot about the corporation's organisation and operations. Here, the competitor would have the same kind of knowledge of what the corporation is doing as the board itself would have, at least with respect to the justification of the cognitive commitment that is embodied in the intention. But the competitor is not taking part in making the decision, so it may not have the same kind of knowledge of the intention itself as the board members have (although I will also deny this in subsection 3.2).

However, a self-knowledge view of singular intentional action is only worth extending to joint intentional action insofar as it provides a plausible account of practical knowledge in the singular case. Here, Velleman provides a detailed alternative evidentialist account of practical knowledge and intentional action. Since this account is the basis for the most detailed extension of the self-knowledge view thus far—as well as one of the best explications of the self-knowledge view of singular intentional action—it is worth going through Velleman's account of how an intention is justified in the individual case in some detail before considering the extension to joint intentional action.

On Velleman's view, an individual intention simply *is* a belief with a particular content and history of formation. At any given moment, an agent will have various desires or preferences pulling her in different directions. According to Velleman, a reflective agent such as an adult human being will in addition to these desires also constantly have a background desire to know what he himself is doing. At least, such self-knowledge will be what Velleman calls a "sub-agential aim" of the agent, in which case it is merely "implicit in some parts of his cognitive architecture" (2015, ch. 2). This background desire or aim is needed to explain how intentions can lead to and guide action, given that intentions are just a kind of belief. But the existence of such a background desire or aim also has some intuitive plausibility: Consider how unsettling it would be to find yourself ignorant of or mistaken about what movements your body was making.

Now, suppose that it was possible for the agent to form a self-fulfilling expectation that represented itself as bringing about the satisfaction of one of those desires that pull him in various directions. Given that the motivational strengths of the other desires are roughly equal, this desire or aim for self-knowledge will motivate the agent to form a self-fulfilling expectation that represents itself as bringing about the satisfaction of one of those other desires. An action that the agent performs as an appropriate result of this will thus satisfy two desires: the desire for the outcome that the action brings about and the desire for self-knowledge. If the purportedly self-fulfilling expectation is indeed fulfilled, then it will be consciously self-fulfilling in that it represents itself as being the cause of its being true and as formed in response to the agent's motivations for action. Velleman's idea is that there is no reason to think that a reflective agent, who desires self-knowledge and knows he tends to do what he expects himself to do, could not rationally form such self-fulfilling expectations. On this point, which is controversial, I think Velleman is right.

To illustrate, suppose I have a desire to continue to write this paragraph, but also a desire to get up and refill my coffee cup. As a reflective agent, I have a desire for self-knowledge, and I know that I tend to do what I expect to do. Relying on this tendency, I form an expectation that this very expectation will cause me to finish writing this paragraph. My desire for self-knowledge gives me an extra motivation to go ahead and do what I expect to do, that is, to finish the paragraph. I will then, let us suppose, be sufficiently motivated to finish the paragraph so that I cause myself to actually finish it. But if I had instead formed a self-fulfilling expectation to refill my coffee cup, then this would have resulted in my getting up and refilling my

empty cup. My desire for self-knowledge would then have helped cause me to have done this. This requires that I know that I have the intention and that I know that I have a general tendency to do what I expect/intend to do. Note that I could not have done whatever I expected, irrespective of what it was that I expected. If I had formed an intention to do something that I was very averse to doing—having a cup of gasoline, say—then this would indeed have made the prospect of drinking gasoline somewhat more attractive to me, but I would nevertheless prefer to do something else—such as having another cup of coffee or finishing the paragraph—even if this would frustrate my expectation of what I was going to do or was doing. There is thus rational pressure on an agent not to expect to do what she doesn't have sufficient prior motivation to do.

The agent's knowledge that she is φ-ing is non-observational on this account because the intention is a prediction and because it is not derived from prior evidence but rather "rests" on the evidence that the agent is both capable of φ-ing and sufficiently motivated to φ for her to actually φ were she to intend to φ. The justification of the intention, on Velleman's view, is hence a matter of straightforward evidential support.

Velleman thinks that both small egalitarian groups and an institution such as an academic department can make decisions and have intentions, but he only gives a detailed account of the former case: a discussion of a dyad that creates a single intention to go for a walk that they jointly hold. I will critically discuss his treatment of this particular case in the next subsection. But a first thought might be that Velleman's evidentialist account of practical knowledge will make the scope of the extension very limited, since a group or organisation would have to instantiate a complex dynamics of deliberation and decision-making involving many beliefs and desires, including the desire for self-knowledge. Perhaps there are some such organisations, but they are the exception rather than the rule.[11] However, Velleman avoids this consequence by arguing that a public representation or representational act, "an utterance, inscription, or depiction of some kind", can literally *be* an intention that is held by several agents as a group given that it can "belong to more than one agent, in some sense of the word 'belong'" (p. 38).

A public representation can in a straightforward sense belong to several individuals without their forming a collective agent in that sense. For example, a to-do list can

---

[11] If a collective agent can have beliefs, desires and intentions of its own, then it is relatively easy to see how at least a "sub-agential aim" for self-knowledge might become implicit in the workings of a collective agent such as a corporation or a political party for example. The need for the organisation to interact with other individual and collective agents and coordinate its activities over time will put pressure on it to only commit to actions that it will indeed carry out, and to only carry out actions that it has committed itself to perform (see List and Pettit 2011, 178; Clark 1994).

belong to a group of friends who are organising a party together. This could also be the case when it comes to a public representation that, according to Velleman's conjecture, is a single token intention. In this case, all Velleman needs is the beliefs and desires of the participants themselves—there is no need for the group as such to have its own beliefs and desires over and above the beliefs and desires of the participants (for discussion, see Bratman 2017, 48-50). At least, this will be the case if the public representation can interact with each participant's desires and beliefs in the right way. To illustrate, suppose that each member of the board desires to know what the corporation he or she is an employee of is doing intentionally. If a public representation among the board members then represents itself as being the cause of it becoming true that the corporation performs a certain action in response to the members' motivations, then this public representation can arguably be an intention that belongs to them all. This sort of account would allow Velleman to extend his self-knowledge view of intention to shared intention in cases of small-scale joint intentional action, such as that of two people going for a walk together, and perhaps also in cases of intentional actions that are performed by organisations that do not themselves aim for self-knowledge. Velleman's account of this is what we turn to in the next subsection. The question remains as to how the members of a group can have non-observational knowledge of a single token intention.

To conclude, the transposition of accounts of practical knowledge such as the ones sketched above to cases of joint intentional action does not introduce any *novel* problems as far as the justification of the cognitive commitment of a shared intention goes.

3.2 Knowledge of the shared intention

The starting point of Velleman's account of how a single token intention can belong to a group is the idea that an intention is a representation that, when effective, causes behaviour by representing itself as causing it. As we have seen, Velleman also thinks that this representation must have a distinct motivational history in order to be an intention. While he never explicitly discusses this part of his theory of intention in his paper, "How to Share an Intention" (1997), his explanation of how to share an intention is clearly supposed to be compatible with this part of the theory.[12] As Velleman recently clarified, he "tried to show that an epistemic conception of intention enables us to explain [...] how intentions can be shared between two agents [...]" (2015, ch. 1).

---

[12] Velleman focuses exclusively on the causal self-reflexivity of intentions, which he takes from Searle's (1980) account of intention and intentional action. The motivational history that he takes to be essential to intention is absent from Searle's account (see Velleman 1997, 36, n. 15 and p. 44, n. 24). Also, Searle does not think that intention is a kind of belief.

Now, Velleman reasons that a public representation or public representational act could be a single intention that is jointly held by several individuals if that representation or act could have the content and functional role of an ordinary intention. Velleman gives two examples of cases where public representations or representational acts might have this content and function: a Philosophy Department's decision to fill a vacancy in a certain way, presumably through public deliberation and discussion (pp. 29–30), and two people exchanging the following speech acts whereby they jointly decide to go for a walk: "I'll go for a walk if you will"; "Then I will."

While Velleman does not explain how a public representation can provide several agents with non-observational and non-inferential knowledge of what they are or will be doing together, the explanation must lie in the sense in which the representation can 'belong' to them as a group. He starts by explaining how a public representation can be an intention that belongs to an individual agent. Velleman asks us to imagine a man who is not sufficiently motivated to get up from his couch to go for a walk and thus cannot rationally form a self-fulfilling expectation that he will go for a walk—a couch potato. However, the couch potato can nevertheless form a rationally self-fulfilling expectation that he is going to declare, "I'm going out for a walk". And he can do this with the *aim* of thereby making a decision aloud to go for a walk. In aiming to do this, he is not settling that he'll go for a walk; settling to settle to take a walk would arguably be to simply settle to take a walk, and this is precisely what he lacks sufficient motivation to do. Rather, what he settles is to make an assertion with the aim or end of settling to take a walk. Making this statement with this aim in mind may make sense because, when the statement has been made, he may end up with an additional motive for taking a walk that he didn't have before (and he knows that the statement will have this effect—indeed, this is why he is making the statement). For example, he may have a desire not to have spoken falsely, and this desire could tip the motivational scale so that he will go for a walk (perhaps there are other desires that could have the same effect). As a result, "[l]ethargy may prevent his initial motives from setting his legs in motion, but it won't prevent them from setting his mouth in motion so as to bring additional motives to bear on his legs" (p. 40). His motives for taking a walk will outweigh his motives against taking a walk.

In these circumstances, the act of uttering "I'm going out for a walk" is according to Velleman a representational act that represents itself as causing the agent to take a walk. Furthermore, this act represents itself as having this effect on the agent by means of the very words spoken by him. The causal effect need not be immediate, but may be mediated by the agent's memory of having earlier decided aloud to take a walk, by his "remaining decided" to take a walk. Now, the couch potato's desire for self-knowledge will ensure that he will indeed go for a walk; otherwise the spoken purportedly self-fulfilling expectation will turn out to be false and his desire for self-knowledge frustrated. Note that the desire not to speak falsely is here not supposed to be analogous to the desire for self-knowledge in the normal case where the agent can directly form an ordinary mental intention to do something. This wouldn't then

be an epistemic conception of (public) intention. The role of the desire not to speak falsely is rather supposed to explain how it is that the utterance can represent itself as causing the agent to take a walk.

Now, let us take a look at Velleman's case involving multiple agents. Suppose that you and I both would like to go for a walk but only in the company of the other. I say to you: "I'll go for a walk if you will". To which you respond: "Then I will." Since each has a desire not to have spoken falsely, each of these conditional statements can represent itself as causing the speaker to go for a walk with the other if the other is willing to do so as well. Each statement represents itself in this way by engaging the speaker's desire for knowledge about what he is doing or what he will do. According to Velleman, each participant's speech act *is* a conditional intention. If all this is common knowledge between us, then our exchange of utterances can jointly represent itself as jointly causing us to go for a joint walk. It does this by tipping the motivational scale in each of us in favour of going for a walk with the other, rather than doing something else. According to Velleman, the upshot of all this is the following:

> Our statements [...] combine to form a joint statement saying, in effect, that they will jointly prompt us to take a walk; and they jointly prompt us to take a walk, as they jointly say. They consequently add up to a single representation that causes our actions by representing itself as causing them—a single token intention that is literally shared between us.
>
> (Velleman 1997, 47)

Is Velleman's characterisation of this joint making up of our minds compatible with his own epistemic conception of intention? According to this conception, an intention provides an agent with non-observational knowledge of what he is doing. Presumably, Velleman's thought is that the single token intention that we share belongs to us in the same way that the spoken intention belongs to the couch potato in the individual case. But the role of the public representational act is very different in the individual and in the joint case. In the individual case, the couch potato is not registering what he is saying as a passive listener. He is the one uttering "I'm going for a walk", after all. Since he is the one who is making the statement—and who intends to make it—he has non-observational knowledge of making it. According to Velleman's account of individual intentional agency, it ought to be in virtue of this non-observational knowledge that the representation belongs to him.

Note that it is not actually necessary for our protagonist to hear his spoken words. For the extra motive to kick in, it is sufficient that he has knowledge that he has spoken them. If he doesn't hear what he himself says because he is listening to loud music with headphones for example, then the extra motive will still kick in. This does not mean that the auditory feedback cannot play an important cognitive role. For instance, hearing the words as he speaks them might make his memory of having spoken them more robustly encoded and easily retrieved. Nevertheless, the hearing of the words does not play any constitutive role in turning his speech act into a spoken intention. Suppose that what was crucial here was that the couch potato did hear the utterance "I'm going for a walk" articulated in his own voice.

Then the utterance could play the same role if it was spoken by someone who had a very similar voice, such as the couch potato's twin brother for example, who was hiding behind the couch. If the couch potato hears the utterance but is unaware of his brother's presence, then he might mistakenly infer that he must have been the one who uttered it, perhaps in order to cause himself to go for a walk. Here, falsely recalling that he uttered these words could have the same causal effect as if he had uttered them himself. The mistaken memory in combination with his desire for self-knowledge could here tip the motivational scale in the same way that an intention normally does. The words spoken by his twin brother would hence cause him to go for a walk. But in this case, the couch potato would not be intentionally going for a walk, at least not in a full-blooded sense.[13] After all, he never settled on going for a walk. Rather, his twin brother made a decision to manipulate him to go for a walk, out of brotherly concern for his health and need for exercise, say. In rising from his couch and starting to walk, the couch potato doesn't have non-observational knowledge of what he is doing under the description "I'm going for a walk". For a speech act to provide the agent with that knowledge, the agent must also be the speaker.

In the joint case, the words spoken play a very different role than in the singular case. It is not sufficient that each of us merely verbally states his or her conditional intention as described above. Each must also recognise and understand the other's statement and believe that it was sincerely made. As Velleman puts it: "if [the] statement 'I will if you will' had been fully precise and explicit, it would have carried an additional condition: 'I will, if you will and if I recognize you as willing'" (p. 46). Hearing the words spoken by the other and being confident that the other has heard oneself doesn't merely play a role in facilitating that the intention is remembered so that it can appropriately cause the represented behaviour. Rather, listening to the other and observing the other listening in turn is constitutive of the representational act of jointly deciding aloud what to do together. This means that the knowledge which each of us has of the other's conditional intention is observational. And so is our knowledge of the other's knowledge of our own conditional intention. Each participant's knowledge of the discharging of the conditional intentions—including his or her own conditional intention—will also be observational. Hence, the knowledge each has of the jointly constituted self-reflexive representation will be observational too. The alleged intention will thus be derived from immediately prior evidence. It will not merely "rest" on evidence once formed.

According to Velleman's own epistemic conception of intention, this means that there is no single intention that we share. In light of this conception, each of us has the wrong epistemic relation to the jointly constituted public representation since each has the role of a spectator with respect to it. It will belong to us in neither the way that the spoken intention belongs to the couch potato, nor in the way that an ordinary non-public intention belongs to an individual agent. The result is that, in

[13] He would not exercise "autonomous agency" (Velleman 1989, 5).

light of the self-knowledge view, we do not have a shared intention. What we have is merely a shared aim with respect to which each would have a spoken conditional intention.

This is not just an objection that is specific to Velleman's particular example of how a shared intention can be created. Insofar as an intention of a group is something that is created publicly among group members—such as the chairman of the board of our corporation uttering "We hereby decide to open the stores", for example—the same problems will arise.

3.3 A potential rejoinder

Philosophers who take intentions to be representations typically take them to be, necessarily, *mental* representations. Intriguingly, Velleman (1997) suggests that his claim that intentions can be oral, written or depicted doesn't necessarily imply a rejection of this assumption. Perhaps publicity and mentality are not mutually exclusive, so that a sound pattern can be the vehicle of some conventional symbols while at the same time be part of the vehicle of some mental content of a person (or persons) who is interacting with that pattern (see Andy Clark 2006, 292–93; Blomberg 2009, sect. 5.1). Now, if we not only share a public representation that appropriately causes us to go for a walk, but share a mental representation that does this, then one could perhaps argue (although Velleman doesn't) that our perception of this representation would be more akin to introspection than to observation of one's environment. Hence, in response to my criticism in subsection 3.2, one could argue that, while the members of a group must indeed rely on perceptual information in order to acquire knowledge of the intention that they share, this acquisition is not really based on *observation* in a sense that conflicts with Velleman's self-knowledge view of intention. Recall that proponents of the self-knowledge view acknowledge that perceptual information plays an important role in enabling the guidance and monitoring of action performance. Perhaps the perception of a public intention that belongs to one's group could be similarly non-observational if the perception were to play an enabling rather than a justifying role.

The idea that the content of representations external to a person's skin and skull can be the content of her mental states is now familiar in the philosophy of mind (see, e.g., Houghton 1997; Clark and Chalmers 1998). In the "The Extended Mind" (1998), Clark and Chalmers argue that a resourceful Alzheimer's patient named Otto has the standing belief that the Museum of Modern Art is on 53rd street, even if the address cannot be retrieved from his biological memory. Instead, he has the belief in virtue of having scribbled down the address in his notebook that he always carries around with him. When he forms the intention to go to the museum, he simply takes the notebook, looks up the address, and proceeds to walk there. From the perspective of commonsense functionalism, the information in the notebook plays the same role as the information that another character, Inga, retrieves directly from biological memory. According to Clark and Chalmers, this should lead us to conclude that a person's standing beliefs can supervene not only on the biological brain and body, but also on various props and tools in an agent's environment.

Clark and Chalmers consider the objection that a crucial difference between Otto and Inga is that Otto has to rely on perception to access the information in the notebook, whereas Inga has more direct access to the content of her beliefs. Because of this difference in epistemic access, the content of Otto's notebook cannot count as the content of his beliefs, the objection goes. But this begs the question, since Clark and Chalmers are arguing that Otto's internal cognitive machinery and the notebook jointly constitute a single cognitive system. Thus, "from the standpoint of this system, the flow of information between notebook and brain is not perceptual at all; it does not involve the impact of something outside the system. It is more akin to information flow within the brain" (1998, 16).[14]

Velleman could perhaps argue that a group member's perception and understanding of a reliable and trusted fellow group member would be somewhat akin to Otto's perception of his notebook. The perception of the intentional actions of in-group members (including their utterances) might best be characterised as a kind of direct assimilation of information, lacking the epistemic vigilance that might be involved in communication with out-group individuals. Indeed, in a discussion of members' identification with the group agent they belong to, List and Pettit (2011) suggest something along these lines.

List and Pettit argue that, just as an individual agent must identify as the agent of its attitudes, so must a group agent identify as the subject of its attitudes. I might believe that a shopper has been making a mess by leaving a trail of sugar behind his cart but, if I fail to realise that I am that shopper, then the belief cannot appropriately figure in an explanation of my behaviour (my apologising to the store clerk, say) (Perry 1979). In normal circumstances, I do not first form a belief about myself in third-person terms and then bridge this identification gap through a further cognitive achievement. Normally, identification is simply, as List and Pettit put it, a "by-product" of how I am constructed as a subject of intentional attitudes. Similarly, members of a well-functioning collective agent automatically identify with this agent's attitudes—"our" attitudes—due to the construction of the collective agent rather than to some further cognitive achievement of each member:

> As members we need not laboriously think of the group in third-personal terms, tracking its attitudes and then working out what is required of us as individual enactors. [...] [T]he members of a group may connect themselves directly to the attitudes of the group. They do not treat the group attitudes as mere indicators of what the group is to do, asking themselves explicitly whether they wish to identify with the group [...]. Rather, their individual attitudes are under the automatic guidance of the group [...].
>
> (List and Pettit 2011, 191–92)

---

[14] An agent does not have standing beliefs in virtue of just any information or resource that is located outside her skin and skull. Certain conditions of "glue and trust" must be satisfied (Clark 2010, 83). For example, the resource must typically be used and invoked by the agent, and the information in it must be easily accessible as well as more or less automatically endorsed.

List and Pettit draw an analogy here to how expert pilots learn to "hitch their intuitions and instincts directly to the instruments" in order to avoid potentially misleading bodily cues that are insensitive to the difference between gravity and acceleration for example (p. 192): "They let their instruments guide them without the intrusion of thoughts about the evidence provided" (*ibid.*).

Unfortunately, none of this will help Velleman or other proponents of the self-knowledge view.[15] With respect to the Extended Mind thesis, the difference between Otto's and Inga's modes of access to the museum's address indeed does not matter to whether they each have a *standing* belief that the museum is on 53rd street. But they both have ordinary privileged access (however that is best understood) to the content of that belief when they are each occurrently thinking about the museum's address (Clark and Chalmers do not suggest otherwise). It is true that Clark and Chalmers also argue that Otto's occurrent sub-personal cognitive processes of memory storage and retrieval could be extended, but this does not mean that his occurrent personal-level mental state of believing that the Museum of Modern Art is on 53rd street would also be extended. But when a group makes a decision and forms an intention, then this cannot be merely a standing intention. At the moment of formation, it must be an occurrent intention, just as Otto's belief that the Museum of Modern Art is on 53rd street must be occurrent as he enters the address in his notebook.

List and Pettit's characterisation of members' behaviour in a well-functioning group agent is not about standing attitudes of the group agent, but about attitudes that are occurrent in the sense that group members ("enactors") are presently activated by them in various ways. However, the by-product model is concerned with how members treat their knowledge that the group they self-identify with has various attitudes *once they know what the content of those attitudes is*. The justification of the members' beliefs that they collectively have attitudes with specific content must at least be partly observational and inferential. There is no way in which group members' knowledge of their intention can be construed as perceptual but non-observational when their intention is a public representation or is formed through some public representational act.

4. Collective joint practical knowledge

At this point, some proponents of the self-knowledge view are likely to object that the extension of the view that I'm criticising is too individualistic. After all, the problems I have raised concern the group members' knowledge of their intention. But what about *the group's* practical knowledge of *its* intention? Or perhaps, what about the members' knowledge of their intention, *qua* group members?

---

[15] To be clear, neither Clark and Chalmers nor List and Pettit are proponents of the self-knowledge view of intentional action.

According to Stoutland (2008), both dyads and other small groups and larger groups or organisations with authority structures and potentially changing membership can be "social agents" that have intentions, perform intentional actions and have practical knowledge of what they are doing. Intentions, beliefs and states of knowledge are according to Stoutland properties that are predicated of a subject that can be a social agent as well as an individual agent (see also Laurence 2011, 292–93). He refers to groups of the former small-scale egalitarian type as "plural agents" and to those of the latter type as "collective agents" (*ibid.*, p. 535). In Stoutland's view, intentions cannot be jointly held or shared in any other way but by being held by a social agent. In particular, he argues that several agents cannot share an intention by each having an intention with the same content since intentions "necessarily include reference to the agent who has them" (Stoutland 2008, 540).[16] Hence, I cannot intend you to make dinner with me, nor can you intend me to make dinner with you. Instead, it is *we*, as a plural agent, who intend to make dinner. In so far as other attitudes are supposed to play a role in explaining the action of a social group, this is also the only way in which they can be shared or jointly held by several agents (*ibid.*, p. 542). Now, in Velleman's account, there is only a 'you' and an 'I' that jointly hold an intention, but no "us" that is a plural agent that holds an intention and performs an action. The knower, in Velleman's account, isn't the doer.

To illustrate the idea that a social agent can have joint practical knowledge, Stoutland provides the following example:

> A corporation, for instance, decides to do something, and then its employees are instructed to carry out the decision; if things do not go as decided, the problem is not that the corporation is wrong about how things went but that the decision was not properly executed, that external conditions changed, or that things went wrong in some other way. The mistake with regard to how things went is not in the judgment (the reports) but in the performance. When it is discovered by investigation that things did not go as intended, there is *theoretical* knowledge of the action. But if things did go as decided, the corporation knows what it is doing simply because its decision was carried out as intended: it has *practical* knowledge of its action, knowledge of what is happening by doing it.
>
> (Stoutland 2008, 546)

The employees are themselves in some sense parts of this collective agent, carrying out the decision of the board, and, as Stoutland himself acknowledges, this board is itself a social agent (*ibid.*, p. 539, n. 7). Interestingly, Stoutland also suggests that we could interpret the following case from Anscombe as an example of a social agent having practical knowledge: "Imagine someone directing a project, like the erection of a building which he cannot see and does not get reports on, purely by giving orders" (1969, sec. 45). Suppose that the man in question is the manager of a team of construction workers, each of whom meticulously and reliably follows the manager's orders. Stoutland's suggestion, then, is that we ought to say that it is the

---

[16] Stoutland is assuming that this reference must be part of an intention's content, but it could also be part of its mode (see Petersson 2015).

team that intentionally erects the building and has practical knowledge that it is doing this.

In both these cases, there is a clear hierarchy with one node at the top, an agent that gives orders to other agents lower in the hierarchy. In Stoutland's own example, this agent is itself a social agent. But insofar as the board is simply treated as a black box, there is nevertheless only one agent at the top of the hierarchy. Laurence (2011, 290–91) uses a similar example to illustrate the possibility of joint practical knowledge. He describes the case of a ship at sea performing a certain action and asks us to imagine that the ship's action is the outcome of the captain giving an order to the first mate, who in light of this order gives several lower-level orders to individual crew members. Again, we find a hierarchy of orders originating from one agent at the top. According to Laurence, we can assemble a practical syllogism from the non-observational knowledge that the captain, the first mate and the crew members each have, even if the captain alone has non-observational knowledge of what the overarching purpose of the ship's action is.

All these cases make it relatively clear how some social agents could have what we might call joint practical knowledge in virtue of some properly authorised individual member's practical knowledge. Giving orders to efficient and trustworthy employees to do something might be a bit like just doing something oneself using one's own hands or tools that one is skilled in using. However, the actions in these cases are at best degenerate cases of joint or shared action (see List and Pettit 2011, 59; Schmid 2016, 65). They are more accurately characterised as individual actions, albeit ones that are performed by appropriately situated individual agents whose possibilities for action have been extended and amplified by specific organisational contexts. Furthermore, the cases appear to be cherry-picked to avoid what is philosophically puzzling about the very idea of joint practical knowledge. They allow the top member's ordinary individual practical knowledge to constitute the practical knowledge that is allegedly also attributable to the whole group or organisation. But Stoutland and Laurence both intend their accounts to also apply to small-scale egalitarian cases of joint intentional action such as that of two people nailing a long board together (Stoutland 2008, 535), or several men pushing a car to a gas station together (Laurence 2011, 286–87). In these cases, there is no single individual agent that can provide the practical knowledge of the whole organisation. Rather, several group members must instead somehow share practical knowledge of one and the same intentional action. What proponents of the self-knowledge view have to do is to explain how this is possible. If we open the black box of the decision-making board from Stoutland's case, we find a group of board members who deliberate together and jointly make decisions on behalf of the corporation. How can they together have joint practical knowledge?

A potential move here would be to argue that it is a mistake to open the black box. The cases in which several agents jointly make a decision could be assimilated to the model of there being one single individual (group-level) agent that uses individual (member-level) agents as social tools. But it is hard to make sense of this claim, unless the group is an *additional* agent with a mind of its own, with the lower-level

individual agents merely having the role of parts in the supervenience base of this additional agent's mind (this is a possibility allowed by, for example, List and Pettit 2011, 66). Again, this would not be a case of *joint* or *shared* agency, and it is clearly a move that no proponent of the self-knowledge view would be willing to make.[17] As Schmid (2016, 66) puts it, "the plural is not just another singular".

As I mentioned at the beginning of section 3, Velleman, Laurence, Rödl, and perhaps also Schmid argue that several agents can literally jointly hold a single intention. But as Laurence makes clear, such a view implies the following:

> *[S]ome* suitably placed person(s) must know without observation what purposes the group is pursuing. A group cannot be said to φ intentionally, if none of its members knows that it is φ-ing, or even if some do but only through observation.
>
> (Laurence 2011, 288)

The question is how some suitably placed *persons*—rather than some suitable placed *person*—can have non-observational knowledge of what the group is doing without this requiring that each of them has non-observational knowledge of the others' non-observational knowledge of what they are jointly intentionally doing. But we saw in the previous section that giving an account of this on the premise that there is a single intention jointly held by several agents runs into what looks like insurmountable problems. At least, such insurmountable problems arise when the intention either is itself, or else is created by, some public representational act.

As far as I can see, what remains for the proponent of the self-knowledge view to do is to dig in his heels and insist that the group's knowledge of its intention can be non-observational even if the members' knowledge is observational. But it is hard to make sense of this claim in light of Laurence's appropriate constraint that "[a] group cannot be said to φ intentionally, if none of its members [non-observationally] knows that it is φ-ing". If such a match in the *content* of knowledge between the group and some group members is required, then a parallel match in the *kind* of knowledge that the group has and that the group members have should also be required.

5. Conclusions

Attempts to extend the self-knowledge view of singular intentional action to joint intentional action face a trilemma: they require that different agents have common non-observational knowledge of what they are jointly intentionally doing, they can only account for the intentional actions of hierarchically structured groups with one agent at the top that is calling the shots, or they will require that the group's intention is the result of a public representational act of which the group members

---

[17] See (Velleman 1997, 38, n. 18); (Stoutland 2008, 543); (Laurence 2011, 291–94); (Rödl 2015, 773); and (Schmid 2016, 66).

will have observational knowledge. The first option would require that the participating agents are not different but are really one and the same agent, or else are for all practical purposes clones of each other with the same preferences, thoughts and reasoning processes. The second option means that the only cases that the extension of the self-knowledge view would be able to account for are cases that at best exemplify limiting cases of joint intentional action. More appropriately, one might think, they should be characterised as cases of socially extended individual action. Finally, the third option implies that the group members will have observational knowledge of the intention that they supposedly jointly hold. For each, the intention will be derived from immediately prior evidence. Hence, it will not be their intention, at least not given the core commitments of the self-knowledge view.

My argument here has focused on cases in which there is some public representation or representational act in virtue of which a group of people jointly makes a decision. I have acknowledged that perhaps there are cases in which a group of people spontaneously acquire a shared intention without there being any such joint decision or agreement. I have not proven that it is impossible that participants could somehow have joint practical knowledge of what it is they are doing jointly intentionally in such a case. But the burden of proof rests squarely on the proponents of the self-knowledge view to show how knowledge of such a spontaneously acquired shared intention could be non-observational. Furthermore, cases in which a group acts on an intention that is formed through a joint decision represent one kind of paradigm case of joint intentional action. An extension of the self-knowledge view ought to be able to capture this kind of case.

Suppose that no extension of the self-knowledge view can avoid the trilemma. What could then be concluded? In section 1, I suggested that focusing on singular intentional action in one's reflection on the nature of human agency is likely to introduce various distortions and biases in one's theorising. Indeed, I think that the idea that practical and non-observational knowledge is constitutive of intentional action is such a distortion.

There is some room for proponents of the self-knowledge view to push back here. They could argue that perhaps singular intentional action and joint intentional action are not subspecies of a general phenomenon of *intentional action*, but rather very different kinds of exercises of human agency. Interesting commonalities could still exist between the singular and the joint cases. Proponents of the self-knowledge view could also point out that there are differences between singular and joint cases that could block extensions of a competing causal control view of intentional action as well, in which case it is at least not so clear that my argument works in favour of an alternative causal control view of intentional action. According to such extensions, some interpersonal pattern of intentions and beliefs is meant to play a role that is analogous to the role that intention plays in singular intentional action (see, e.g., Tuomela and Miller 1988; Bratman 1992, 2014; Pettit and Schweikard 2006; Tuomela 2007; Ludwig 2016). Many proponents of such a causal control view would accept that intentions are essentially indexical mental states, but there is no

immunity to error through self-identification when it comes to the question of whether you and I instantiate an interpersonal pattern of intentions and beliefs that can effectively control and coordinate our joint action (Searle 1990, 408; Bratman 2014, 58; Schmid 2016, sect. 2). I can think that such a pattern is in place, but I might be mistaken and you have no part in it. Does this imply that, given the causal control view, intention and intentional action cannot be essentially indexical after all? Or that proponents of the causal control view should be eliminativists about shared intention and joint intentional action? Or does it show that singular intentional action and joint intentional action are actually very different creatures? I leave these questions for another occasion.[18]

References

Anscombe, G. E. M. 1969. *Intention*. Ithaca: Cornell University Press.

Blomberg, Olle. 2009. "Do Socio-Technical Systems Cognise?" In *Proceedings of the 2nd Aisb Symposium on Computing and Philosophy*, 3–9.

———. 2011. "Socially Extended Intentions-in-Action." *Review of Philosophy and Psychology* 2 (2): 335–53.

———. 2016. "Common Knowledge and Reductionism about Shared Agency." *Australasian Journal of Philosophy* 94 (2): 315–26.

Bratman, Michael. 1992. "Shared Cooperative Activity." *The Philosophical Review* 101 (2): 327–41.

---

[18] For comments, questions and suggestions, thanks to Facundo Alonso, Michael Bratman, Adam Carter, Herbert Clark, Nathan Hauthaler, Carlos Nunez, Orestis Palermos, Grace Paterson, Herlinde Pauer-Studer, Björn Petersson, Alessandro Salice, Glenda Satne, Hans Bernard Schmid, Will Small, Alexander Stathopoulos, and John Turman. Thanks also to the participants at the Berkeley Social Ontology Group meeting on April 20, 2016, the Varieties of Agency workshop at Stanford on May 11, and the Collective Intentionality X conference in Den Haag on September 1, where ideas in this chapter were presented. Most of the work on this chapter was carried out while I was a Visiting Scholar at the Department of Philosophy at Stanford University in the Winter and Spring of 2016. I am very grateful to Michael Bratman and the Sweden-America Foundation for making this extended visit possible. The research was funded by project grant 421–2014–1025 from Vetenskapsrådet (the Swedish Research Council).

———. 1997. “I Intend That We J.” In *Contemporary Action Theory: Social Action*, edited by G. Holmström-Hintikka and R Tuomela, 49–63. Dordrecht: Kluwer.

———. 2009. “Intention, Belief, Practical, Theoretical.” In *Spheres of Reason: New Essays on the Philosophy of Normativity*, edited by Simon Robertson, 29–61. Oxford University Press.

———. 2014. *Shared Agency: A Planning Theory of Acting Together*. Oxford University Press.

Bratman, Michael E. 2017. “The Intentions of a Group.” In *The Moral Responsibility of Firms*, edited by Eric W. Orts and N. Craig Smith, 36–52. Oxford University Press.

Clark, Andy. 2006. “Material Symbols.” *Philosophical Psychology* 19 (3): 291—307.

———. 2010. “Memento’s Revenge: Objections and Replies to the Extended Mind.” In *The Extended Mind*, edited by Richard Menary, 43–66. MIT Press.

Clark, Andy, and David Chalmers. 1998. “The Extended Mind.” *Analysis* 58 (1): 7–19.

Clark, Austen. 1994. “Beliefs and Desires Incorporated.” *The Journal of Philosophy* 91 (8): 404–25.

Clark, Herbert H. 2006. “Social actions, social commitments.” In *Roots of Human Sociality: Culture, Cognition and Human Interaction*, edited by Stephen C Levinson and N J Enfield, 126–50. Oxford: Berg Press.

Falvey, Kevin. 2000. “Knowledge in Intention.” *Philosophical Studies* 99 (1): 21–44.

Grünbaum, T. 2009. “Anscombe and practical knowledge of what is happening.” *Grazer Philosophische Studien* 78 (1): 41–67.

Houghton, David. 1997. “Mental Content and External Representations.” *The Philosophical Quarterly* 47 (187): 159–77.

Laurence, Ben. 2011. “An Anscombian Approach to Collective Action.” In *Essays on Anscombe’s Intention*, edited by Anton Ford, Jennifer Hornsby, and Frederick Stoutland, 270–96. Harvard University Press.

List, Christian, and Philip Pettit. 2011. *Group Agency: The Possibility, Design and Status of Corporate Agents*. Oxford University Press.

Ludwig, K. 2016. *From Individual to Plural Agency: Collective Action (Volume 1)*. Oxford, UK: Oxford University Press.

O’Brien, Lucy F. 2003. “On Knowing One’s Own Actions.” In *Agency and Self-Awareness: Issues in Philosophy and Psychology*, edited by Johannes Roessler and Naomi Eilan. Clarendon Press.

Perry, John. 1979. “The Problem of the Essential Indexical.” *Nous* 13 (1): 3–21.

Petersson, Björn. 2015. “Bratman, Searle, and Simplicity. A comment on Bratman: Shared Agency, A Planning Theory of Acting Together.” *Journal of Social Ontology* 1 (1): 27–37.

Pettit, Philip, and David Schweikard. 2006. “Joint Actions and Group Agents.” *Philosophy of the Social Sciences* 36 (1): 18–39.

Rovane, Carole. 1998. *The Bounds of Agency: An Essay in Revisionary Metaphysics*. Princeton University Press.

Rödl, Sebastian. 2007. *Self-Consciousness*. Cambridge, Massachusetts: Harvard University Press.

———. 2015. “Joint action and recursive consciousness of consciousness.” *Phenomenology and the Cognitive Sciences* 14 (4): 769–79.

Schmid, Hans Bernhard. 2016. “On Knowing What We’re Doing Together: Groundless Group Self-Knowledge and Plural Self-Blindness.” In *The Epistemic Life of Groups*, edited by Michael Brady and Miranda Fricker, 1–31. Oxford University Press.

Schuetz, Alfred. 1953. “Common-Sense and Scientific Interpretation of Human Action.” *Philosophy and Phenomenological Research* 14 (1): 1–38.

Searle, John. 1980. “The intentionality of intention and action.” *Cognitive Science* 4 (1): 47–70.

Searle, John R. 1990. “Collective Intentions and Actions.” In *Intentions in Communication*, edited by Philip R. Cohen, Jerry Morgan, and Martha E. Pollack, 401–15. MIT Press.

Seemann, A. 2009. “Why We Did It: An Anscombian Account of Collective Action.” *International Journal of Philosophical Studies* 17 (5): 637–55.

Setiya, Kieran. 2008. “Practical Knowledge.” *Ethics* 118 (3): 388–409.

———. 2009. “Practical Knowledge Revisited*.” *Ethics* 120 (1): 128–37.

Small, W. 2012. “Practical Knowledge and the Structure of Action.” In *Rethinking Epistemology 2*, edited by Günter Abel and James Conant, 133–227. Walter de Gruyter.

Sosa, Ernest. 2015. *Judgment and Agency*. Oxford University Press.

Stathopoulos, Alexander. 2016. “Knowing Achievements.” *Philosophy*, April, 1–14.

Stoutland, Frederick. 2008. “The ontology of social agency.” *Analyse & Kritik*, 533–51.

Thompson, Michael. 2008. *Life and Action: Elementary Structures of Practice and Practical Thought*. Harvard University Press.

Tuomela, Raimo. 2007. *The Philosophy of Sociality: The Shared Point of View*. OUP USA.

Tuomela, Raimo, and Kaarlo Miller. 1988. "We-Intentions." *Philosophical Studies* 53 (3): 367–89.

Velleman, J David. 2007. "What good is a will?" In *Action in Context*, edited by Anton Leist, 193–215. Berlin & New York: de Gruyter.

Velleman, J. David. 1989. *Practical Reflection*. Princeton University Press.

———. 1997. "How to Share an Intention." *Philosophy and Phenomenological Research* 57 (1): 29–50.

———. 2015. *The Possibility of Practical Reason*. 2nd ed. Michigan Publishing.