

Article

Open Access

Sveinung Sundfør Sivertsen

No Need for Infinite Iteration

A Critique of the Collectivist Copernican Revolution in Social Ontology

DOI 10.1515/jso-2014-0026

Abstract: As part of his argument for a “Copernican revolution” in social ontology, Hans Bernhard Schmid (2005) argues that the individualistic approach to social ontology is critically flawed. This article rebuts his claim that the notion of mutual belief necessarily entails infinite iteration of beliefs about the intentions of others, and argues that collective action can arise from individual contributions without such iteration. What matters is whether or when there are grounds for belief, and while extant groups and social structures may be relevant to some forms of collective action, this does not show that all forms of collective action depends on such such pre-established collectivity.

Keywords: Collective action; Individualistic approach; Mutual belief; Infinite iteration; Common knowledge.

1 Introduction

In part two of his book, *Wir-Intentionalität* (2005), Hans Bernhard Schmid argues for a “Copernican revolution”¹ in social ontology, through which the “we” of collective humanity will come to be considered ontologically primary to the “I” of the individual. Schmid’s exploration of when and how we come to act and think

¹ “ii. Eine Kopernikanische Wende” (Schmid 2005, p. 107). This and other translations from German by the author.

Sveinung Sundfør Sivertsen, Department of Philosophy, University of Bergen, Bergen, Hordaland, Norway, e-mail: sveinung.sivertsen@uib.no; sveinung.sundfor@outlook.com

of ourselves as “we” is lucid and engaging, but his criticism of the individualistic approach to social ontology rests on an error, or so it will be argued here.

First associated with Max Weber (1922) and more recently championed by, amongst others, Raimo Tuomela (Tuomela and Miller 1988; Tuomela 2005), the individualistic approach argues that individual agents form teams, groups or collectives through appropriate combinations of individual intentions, behaviour, and, usually, some form of mutual belief. So, for example, if you and I were to carry a table together, we would have to behave appropriately, namely by lifting opposing ends of the table; we would have to intend our lifting as contributions to the collective act of carrying the table together; and we would have to have the belief, mutual between us, that each of us would be doing our part contributing to the collective act. If any of these were to be missing, we would not perform the collective act, that is, we would not be lifting the table. The exact details of what goes into the forming of groups varies according to the flavour of the individualistic thesis, but here I will simply rely on Schmid’s own characterisation of this kind of approach.

In a section of the book entitled “Reductionism: a Dead End” (Schmid 2005, p. 108–147),² Schmid argues that the individualistic approach fails for two reasons: 1) because the kind of belief it requires for an act to count as a contribution to a collective act entails an infinite iteration of “A thinks that B thinks that A...”, etc., and 2) because even if one were to allow such psychologically improbable beliefs, this still does not beget the mutual belief or common knowledge required for an act to count as a contribution to a collective act.³ This article aims to show that Schmid is wrong in both these claims.

2 The Third Criterion

In order to discuss the individualistic approach, Schmid introduces an example of two ramblers – Anna and Bertha. After a chance encounter walking the same mountain trail, the pair play out variations of coordinated action, while Schmid considers the various thoughts, beliefs and intentions that each of them could have, and what these would do to influence the degree of collectivity in their

² “§4 Reduktionismus: eine Sackgasse” (Schmid 2005, p. 108)

³ Schmid appears to treat mutual belief and common knowledge as interchangeable, and, for the purposes of the present paper, I shall do the same, ignoring any difference there might be in the rational demands of knowledge compared to those of belief.

actions (2005, p. 107–132). By way of examples, Schmid quickly arrives at the conclusion that social coordination on its own does not beget collective action. If Anna, who is walking ahead of Bertha, simply reduces her speed, this does not count as a contribution to group formation; she might simply be tired.⁴ Even if Anna slows down *in order to* wait for Bertha, this still does not count – she might just want Bertha to pass her by, so that she yet again can walk undisturbed on the path. What is required is that Anna slows down in order to wait for Bertha *so that* they can continue walking together as a group. Anna must intend her behaviour – slowing down – as a contribution to the formation of a walking group where Bertha is the other party.

However, if Anna is alone in intending her coordinated behaviour – slowing down – as a contribution to the formation of a walking group, and Bertha, on her part, intends her coordinated behaviour – speeding up – simply as a way of passing Anna by, this clearly does not result in a walking group (but rather a case of awkward social coordination). And so, in order for there to be any chance of Anna and Bertha forming a walking group, they must intend their respective behaviour as contributions to the collective act of walking together.

The first criterion of collective action is therefore that each would-be member of a group act in an appropriately coordinated manner, while the second criterion is that each would-be member intends this behaviour as their contribution to the collective act.

But, continues Schmid, Anna could not form the intention to walk together with Bertha if she thought that Bertha would not walk together with her. Thus, it seems that Anna, in order to form the intention requisite for her socially coordinated behaviour to count as a contribution to the collective act, must expect or believe that Bertha will in fact walk together with her. Now, Schmid says here that Anna must believe that Bertha *intends* to contribute to the collective act (Schmid 2005, p. 113–114). There is, I think, an equivocation in ordinary language that comes into play at this point. If I expect that you will do something, in this case that you will do your part contributing to our collective endeavour, does this necessarily mean that I form an explicit belief that you have an intention of so doing? Or does it suffice that I form a belief or an expectation that you will do so? I have my doubts about the need for ordinary thinking about the actions of others to include explicit beliefs about their intentions. However, since it is almost universally accepted that ordinary thinking about the potential contributions of others to collective acts necessarily involves forming explicit belief about

⁴ For what I have termed “contributions,” Schmid uses the terms “Beitragshandeln/-verhalten” – literally contributive action/behaviour.

their intentions of so doing, I shall leave this doubt aside here.⁵ Accepting that beliefs about intentions are required, we can go along with Schmid and propose a formulation for a third criterion of collective action: Each would-be member of a group must believe that the other would-be members have the intention to contribute to the collective act. We now have the following three criteria for deciding whether a collective act occurs:

- C1 Each would-be member of a group acts in an appropriately coordinated manner
- C2 Each would-be member intends their behaviour as their contribution to the collective act
- C3 Each would-be member believes that the other would-be members have the intention to contribute to the collective act⁶

Note that these criteria leave room for failed individual contributions to be recognised as “would-be contributions” so long as they fail due to the actions or inactivity of others. If an agent believes that the other would-be members have the intention to contribute, intends her behaviour as a contribution, and begins to act in such a way that her action will be appropriately coordinated if the other would-be members contribute, her action will only count as a full-blown contribution if the other would-be members likewise contribute. But even if they do not, her action should count as a would-be contribution, since the only thing on which it stumbles is the lack of such contributions from others. This interpretation of the individualistic approach is supported by Schmid’s (implicit) acknowledgement that there is something like a would-be contribution that fails due solely to the actions or inaction of another (Schmid 2005, p. 110–111). Furthermore, if we were to make the actuality of a collective act a success condition of individual contributions, it seems we would exclude an individualistic approach by definition, since the success conditions for each individual contribution would then have to include the physical acts of others, and not simply require the individual to have *beliefs* about these. Allowing that an individualistic approach should not

⁵ Christopher Kutz (2000) is a notable exception to this rule. His positive argument for a minimal account of joint action where the only base requirement is that agents act with participatory intentions fits nicely with my negative argument against Schmid.

⁶ Schmid’s formulation of the two first criteria differ somewhat from my own here because he is talking of a group consisting of two agents, A(nna) and B(ertha). The third criterion as I have formulated it here does not figure explicitly in Schmid’s discussion, but the requirement that it embodies does. Translated from German, Schmid’s formulation of the first two criteria are: “1) The behaviour of A and B correspond in an appropriate manner. 2) A and B must intend their behaviour as contributions to a collective action.” (Schmid 2005, p. 121).

be excluded by definition, we should therefore allow that individual would-be contributions might count as such even if the collective act does not take place.

Putting failed contributions aside for now, the above formulation of the third criterion is where the individualistic approach gets into trouble according to Schmid. By making the contribution of each would-be member depend on that member's belief about the intentions of the other would-be members, we have, he claims, initiated an infinite iteration of beliefs about beliefs. To see this, consider Anna and Bertha again: In order for Anna's socially coordinated behaviour (C1) to count as a contribution to forming a walking group with Bertha, Anna must intend her behaviour as such (C2). But Anna will not form the intention to contribute unless she believes that Bertha will also contribute, and so Anna must believe that Bertha intends to contribute (C3). So far, so good. But since Bertha will not form *her* intention to contribute unless she believes that Anna intends to contribute, it seems that Anna must also believe that Bertha believes that Anna intends to contribute in order for Anna to form the intention to do so. And since Anna's intention to contribute depends on her belief about Bertha's intention, this seems to imply that Anna must believe that Bertha believes that Anna believes that Bertha intends to contribute. And so on, apparently in an infinite iteration of beliefs.

If Schmid is right, we have to revise our formulation of the third criterion. The revised version would look something like this:

C3* Each would-be member has infinitely iterated beliefs to the effect that the other would-be members have the intention to contribute to the collective act⁷

There are, in Schmid's view, two distinct ways in which this infinite iteration of beliefs creates trouble for the individualistic approach. The first is that it makes an individualistic analysis of collective action seem psychologically implausible. No normal human being can form an infinite number of beliefs in this manner. While there have been attempts at defusing this threat of implausibility,⁸ Schmid rejects these as psychologically implausible for other reasons (Schmid 2005, p. 130–131). The second way in which the revised third criterion creates trouble for the individualistic approach is more serious, according to Schmid. For even if

⁷ Again, my formulation differs from Schmid's, but I believe C3* captures the core of his argument at this particular point in the discussion. Translated from German, Schmid's formulation of the third criterion is: "A and B have an infinitely iterated mutual knowledge of their intention to contribute" (Schmid 2005, p. 121).

⁸ Although Margaret Gilbert's approach is not individualistic, Schmid discusses her "smooth reasoner counterpart" in particular (see Gilbert 1989).

we allow that normal agents can form any number of such iterated beliefs, this still does not appear to give them *rational grounds* for contributing to the collective act. Briefly, there are situations in which would-be contributors to a collective act are unable to establish the requisite mutual belief that the others intend to contribute, even if we allow them to iterate infinitely their beliefs about each other's intentions. I will deal with each of these worries in turn.

Instead of trying to come up with another “defusing strategy”⁹ to take the edge off the threat of psychological implausibility, I will argue that a proponent of the individualistic approach can accept the first formulation of the third criterion (C3) without thereby committing to the revised version (C3*). An individualistic approach need not occasion *any* iteration of beliefs, let alone an infinite one. This will of course only get me so far, since the problem of procuring rational grounds for contributing to collective acts remains unsolved by this conclusion. However, by freeing the individualistic approach from the apparent *necessity* of iterations of belief, there arises another possibility for understanding how individuals can come to have rational grounds for contributing to collective action: one that can include, but need not depend on iterations of belief.

3 No Infinite Iteration

Schmid takes one crucial false step in his argument that C3 implies C3*. Given C2 and C3 for the case of Anna and Bertha, we have the following: Anna must intend to contribute (C2 for Anna), and in order to do so, she must believe that Bertha likewise intends to contribute (C3 for Anna). Bertha must herself intend to contribute (C2 for Bertha), and in order to do so, she must believe that Anna also intends to contribute (C3 for Bertha). From this, Schmid concludes that Anna must believe that Bertha believes that Anna intends to contribute, and so on (C3* for Anna). To arrive at this conclusion, Schmid must treat the two instances of Bertha's intention in the above description – one as had by Bertha (C2), the other as the object of Anna's belief (C3) – as interchangeable. It is the failure to distinguish between these two instances of the intention that sets off the infinite iteration of beliefs.

The distinction might not be immediately apparent, but it is important. In terms of what we would require in order to ascribe beliefs and intentions to people – call these existence conditions –, there are different requirements for a person to count as having an intention and for another person to count as having a belief

⁹ “Entdramatisierungsstrategie” (Schmid 2005, p. 128)

about that intention. To borrow a point from Tuomela, if Bertha really intends to P she must “at least lack the belief that it is impossible for [her] to perform the action” (Tuomela 2005, p. 329). At the same time, it is clear that Anna can both believe that Bertha intends to X *and* believe that it is impossible for Bertha to X. And so the existence conditions for the intention-as-had do not simply transfer to the intention-as-object-of-belief. As a result, Schmid cannot assume that existence conditions pertaining to Bertha having the intention to contribute are shared by Anna having a belief about that intention. That is, Schmid cannot assume that Bertha’s intention to contribute and Anna’s belief about Bertha’s intention both require Bertha’s belief about Anna’s intention to contribute.

In the standard case, then, Bertha’s intention as an object of Anna’s belief cannot be conflated with the intention as had by Bertha, since the existence conditions for each may differ.¹⁰ However, even if one accepts this claim for the standard case, one might object that it does not hold for the special case of collective action. For, if Anna believes that Bertha cannot contribute to the collective act of which Anna is the only other potential member, it would be wildly irrational for Anna to form the intention to contribute in spite of this. The existence condition is shared, after all: Anna must at least lack the belief that it is impossible for Bertha to contribute in order for Anna to form the belief that Bertha intends to contribute.¹¹

Does this mean that the distinction between the two mental objects is inconsequential in the special case of collective action? No. The point of the distinction was to show that an intention as had by a person and that intention as the object of the belief of another person do not *necessarily* share existence conditions. The objection has proved that they share one existence condition, but it remains an open question whether the other existence conditions pertaining to Bertha’s intention to contribute are also shared by Anna’s belief about that intention. Specifically, it remains an open question whether Bertha’s intention to contribute and Anna’s belief about Bertha’s intention both require Bertha’s belief about

¹⁰ It has been suggested to me that a disjunctivist about belief might object to this claim. A disjunctivist would say that the object of Anna’s belief about Bertha’s intention to contribute *just is* Bertha’s intention, and so there is no problematic conflation in play. However, the disjunctivist explains the “unhappy” case of false belief, i.e. where the object of belief is missing, as a case where belief is only seemingly about that object. If we agree that would-be contributions to collective acts should count as such, it is clear that we can require no more of individual agents than their beliefs seemingly being about the intentions of other agents to contribute. My claim that the conflation is illicit is therefore compatible with a disjunctivist view of beliefs, at least for the case of contributions to collective action.

¹¹ Thanks to an anonymous reviewer for providing this objection.

Anna's intention to contribute. For the iteration to get off the ground, Schmid must show that the two mental objects do in fact share *this* existence condition, and, as far as I am concerned, he has not done so.

Here one might attempt an objection to my reply that parallels the one above, and say that it would be wildly irrational for Anna to form the belief that Bertha intends to contribute if Anna also believes that Bertha does not believe that Anna intends to contribute. If this is the case, Anna will not form the belief that Bertha will contribute, and so will not form the intention to contribute herself. In other words, Anna must at least *lack* the belief that Bertha *does not* believe that Anna intends to contribute in order for Anna to form the intention to contribute herself. But this does not help Schmid: saying that Anna must lack this belief is not the same as saying that she must have an explicit belief about what Bertha in fact believes about Anna's intention. Lacking a belief about something being the case is not the same as having a belief about its opposite being the case.

In sum, one can accept the original formulation of the third criterion (C3) without thereby committing to the revised version (C3*). Schmid cannot complain that an individualistic approach gives rise to an infinite iteration of beliefs based on the requirement that would-be contributors to a collective act form beliefs about the intentions of others to contribute. That said, there might be other ways in which an individualistic account invites infinite iteration of beliefs. I look at some of these next.

3.1 Dispositional Knowledge, Ordinary Beliefs

Even if one thinks – for whatever reason – that would-be members of a group need to form beliefs about the *beliefs* of other would-be members, and not simply about their intentions, this still would not lead to an infinite iteration of beliefs, at least not of any kind that renders the individualistic approach psychologically implausible. Not even requiring that agents iterate their belief to some arbitrarily high degree would have this result.

To see this, consider that any kind of belief implies a potentially infinite number of other beliefs, without this stopping us from ascribing beliefs to ordinary agents in other settings. A person holding the belief that Oslo is in Norway need not simultaneously hold the belief that Oslo is in Norway *or on Mars*, even if that is implied by the first belief. The usual epistemological response to the problem of infinite iteration of beliefs, which Schmid somewhat unsatisfactorily treats in a footnote to his discussion (2005, p. 125, footnote 19), is to say that any belief implied by those we hold explicitly can be seen as both dispositional and

implicit, and hence do not encumber the person to whom we ascribe the explicit belief (Schwitzgebel 2014).

Only if the beliefs in question are *explicit* beliefs, i.e. beliefs actually entertained by Anna at a given point in time, can there be any psychological problem. It is not sufficient for Schmid's argument that Anna *could* form these beliefs if she were to pursue a chain of reasoning; the beliefs cannot merely be implicit in her other beliefs, be derivable from them. The simple fact that Anna, reflecting on her belief about Bertha's beliefs about Anna's intentions, is disposed to form additional beliefs of the "I believe that Bertha believes that I believe that Bertha believes" – kind does not in itself pose any problem for the psychological plausibility of the individualistic approach (see Lewis 1969).

3.2 Objection: Not Ordinary Beliefs, Rational Grounds for Action

To this, a defender of Schmid's position might respond that even if it is true that implied beliefs may safely rest undisturbed in the implicit dispositional domain in the case of *most* beliefs, this is not the case for beliefs about the intentions of others to contribute to collective acts. Borrowing a point from Jane Heal (1978), Schmid points out that an adequate conception of mutual belief or common knowledge must not only give us something that is psychologically plausible (i.e. not entailing an infinite iteration of explicit beliefs), but also something that (together with the appropriate desires) provides rational grounds for action. And, continues Schmid, rational grounds for action cannot contain contradictory beliefs (Schmid 2005, p. 126–127).

Given the right circumstances, whether or not an agent has rational grounds for intending her behaviour as a contribution to a collective act may depend on a high order of iterated belief about the beliefs of other potential contributors. Schmid gives the following example: Anna has stopped some ways in front of Bertha on the trail, and Bertha sees this. But when Anna looks back toward Bertha behind her on the trail, Anna is almost blinded by the setting sun. Anna can still make out Bertha against the strong backlight, but she is not sure that Bertha knows that she can see her. Anna then remembers telling Bertha earlier about the way her glasses makes it almost impossible for her to see in backlit situations, and so Anna concludes that Bertha, being intelligent enough to take this information into account, has concluded that Anna cannot see her, and thus cannot see that she sees that Anna has stopped to wait for her. Anna then concludes that Bertha (falsely) believes that Anna does not know that Bertha knows that Anna has stopped to wait for Bertha (Schmid 2005, p. 128–129). In this case,

Anna's explicit, third-order iterated belief about the beliefs of Bertha contradicts the non-iterated belief that Bertha intends to contribute, and so Anna does not have rational grounds for contributing.¹²

The problem is that, given less "intelligent" agents, a situation in which there are grounds for forming beliefs that are contradictory to the ones currently held, but where neither agent follows the iteration up to the level of contradiction, the success conditions of each agent's contribution would be fulfilled. In this case, there would be mutual belief or common knowledge, and the collective act would occur, even if there (supposedly) were no rational grounds for it. I am not entirely sure why this is a problem for the individualistic approach. Schmid himself here asks, rhetorically: what should we say about a situation in which there is common knowledge, even if, following closer examination, rational agents would *not* form common knowledge (Schmid 2005, p. 130)? What indeed? Saying that this should *not* count as common knowledge is simply affirming the consequent, and even if we think there is something less than fully rational in each of the two agents forming the requisite beliefs, this in itself does not pose a conceptual problem for the individualistic approach. Such an approach need not exclusively account for the collective actions of ideally rational agents. As long as it can explain when and how it *would* be rational for an individual to contribute to a collective act, it might well be said to be a strength that a theory is also able to explain collective action by less than fully rational agents.

As for where to set the bar for "rational," we would certainly think it strange if an agent, after having iterated her beliefs to some level and having found a contradiction, then proceeded to act on this contradictory set of beliefs. But it is far less obvious that there is something wrong with an agent who does not iterate her belief, and acts on what she takes to be true, even if she would not take this to be true had she iterated her beliefs. Neglecting to iterate belief in itself does not give rise to any contradictions, and if the charge is that an agent would be irrational to act on a contradictory set of *explicit* beliefs, then the lazy iterator is

¹² Schmid's discussion here differs somewhat from my own, because he is treating the "defusing strategy" of only requiring of agents that they iterate their beliefs to a low degree, e.g. to the third degree, in order for mutual belief to count as common knowledge. He says of this example that, as a consequence of Anna's iteration of beliefs, it is not mutual belief between Anna and Bertha that Anna has stopped to wait for Bertha. Since I do not accept as a default that individuals must iterate their beliefs in order to have rational grounds for contributing to a collective act, I focus here instead on the fact that Anna's iteration leads to her not having rational grounds for contributing. The problem identified by Heal/Schmid, as well as my way of responding to it, apply both to the three-iterations and the no-iterations views, since the source of the problem is not in any number of iterations of belief, but in the grounds that are present for forming a higher-order iterated belief that contradicts a lower-order one.

not being irrational. If the charge is that it is irrational to act on any set of beliefs that contain contradictions, whether the beliefs in question are explicitly entertained or merely implicit and dispositional, the criterion of rationality applied would condemn as irrational much of what we normally call rational behaviour – anyone might discover that some of their beliefs contradict each other. Philosophical inquiry is, since Socrates, to a large extent an exploration such beliefs. To require of us that we be *certain* that we always act on a non-contradictory set of beliefs for our actions to count as rational is clearly too strict.

There might of course be reasons for which we would require that agents iterate their beliefs about other would-be contributors to *some* degree above zero, at least for some kinds of collective action, such as when we are planning a collective endeavour.¹³ For other, simpler cases, I think we should be careful in calling the lazy iterator *irrational*. We could perhaps say that it is epistemically irresponsible not to iterate belief to some level before acting, but I am sceptical of discounting, *a priori*, the rationality of acting on non-iterated beliefs. Even if there are grounds for forming a contradicting belief at a low level of iteration, if neither agent iterates their belief, but both simply assume that the other will contribute and proceed to contribute themselves, they will in fact be acting on correct mutual belief. In such cases, rather than facilitating collective action, iteration of beliefs gets in the way of collective action.¹⁴ One might think that rational agents could recognise this possibility, and adapt a strategy of not overthinking their contributions to collective acts as a way of maximising their successes in such endeavours. If so, non-iteration could sometimes be the rational thing to do.

4 The Problem of Mutual Belief

My argument so far has been negative, arguing against Schmid that an individualistic analysis of collective action need not invoke an infinite iteration of beliefs. This leaves untouched the second problem identified by Schmid, namely that the individualistic approach appears unequipped to explain when would-be contributors to a collective act actually have grounds for forming the requisite beliefs about each other's intentions to contribute. Schmid frames this problem in terms of the inability even of an infinite iteration of beliefs to generate mutual belief

¹³ Thanks to an anonymous reviewer for encouraging this point. See also (Kutz 2000) for examples and discussion.

¹⁴ See Chant and Ernst (2008) for a different way in which iterations of belief can block collective action.

or common knowledge. In keeping with the structure of the argument here, we can think of Schmid's point as a challenge to proponents of the individualistic approach to provide an explanation for how individual agents can come to have rational grounds for their contributions to collective acts. If (infinite) iteration of beliefs is neither necessary nor sufficient for establishing such grounds, what is?

Providing a comprehensive answer to this question is outside the scope of this paper. But I will hint at what shape I think the answer will have to take. In order to do so, I will first look at Schmid's argument for his conclusion about infinite iterations of belief being insufficient to establish mutual belief. Granting Schmid's claim for the special case he uses as an example, I will instead argue that this example is so artificial that the cost of accepting it is vanishingly low. Indeed, the cost is one that any proponent of the individualistic approach should be happy to pay, because would-be contributors to collective acts will seldom, if ever, rely solely on iterations of belief to provide rational grounds for their contributions.

4.1 An Intentional Somersault

After having discussed the problem of infinite iteration, Schmid undertakes an outing into "early phenomenology" in order to provide alternate ways of approaching the problem of collectivity (Schmid 2005, p. 132–141). The gist of this section of Schmid's book is the claim that a collective precedes rather than is created by individually held beliefs in the collective *experience* of something, a move that Schmid, borrowing from Walther, calls an "intentional somersault" (*intentionalen Purzelbaum*). As interesting a read as this is in itself, it depends for its pertinence to the current argument on a problematic assumption, namely that there is a defensible equality to the *prima facie* different phenomena of collective experience and collective action. If Schmid is to derive support from these sources, he needs to show that their insights are relevant to collective action.

Schmid attempts this transference using a pair of complementary examples to show (a) that common knowledge is what renders the (infinite) iteration of individual belief possible and (b) that (infinite) iteration of individual belief cannot produce common knowledge (Schmid 2005, p. 141–144). Together this sets the stage for an "intentional somersault" like that of Walther, whereby commonality of knowledge is put prior to the beliefs individuals might have about the subject of that knowledge. This in turn is the first part of Schmid's proposed Copernican revolution, where commonality of knowledge is explained by the ontological priority of collectivity over individuality. The first example is a version of "The Barbecue Problem" (Vanderschraaf and Sillari 2014), and will not be treated here. I will simply grant Schmid that individual knowledge can be derived from collective

knowledge. The crucial point Schmid needs to prove is that the common knowledge or mutual belief necessary for collective action cannot be constructed from (iterations of) individual belief, and for this, he needs the second example.

The second example is a version of the Coordinated Attack problem, first described in the distributed systems literature (Chant and Ernst 2008). In Schmid's telling of it, two infantry detachments have separated as part of a tactical manoeuvre. They are trying to fix the time of the upcoming attack to six o'clock the next morning, and the first detachment sends a messenger with that message to the second. Given the possibility that the first messenger could have been intercepted, and seeing as how the first detachment cannot know whether the second has received the order or not, the second detachment decides to send a messenger in return confirming the time of attack, thereby assuring the first detachment that they know the order and will commence the attack at the time set. The second messenger also arrives, but given the possibility that this messenger could also have been intercepted, the first detachment reasons that the second detachment cannot know whether the first detachment knows that the second detachment has received the order and will commence the attack at the time set. If they do not know this, the whole structure collapses, making any previous communication worthless in the process. The first detachment therefore sends out a third messenger, confirming that they have received the second detachment's confirmation – *et cetera ad infinitum*. This shows that no amount of back-and-forth can produce the knowledge the two detachments need as rational grounds for commencing the attack at the time set. The exchange of messengers – like the iteration of beliefs – fails to establish the common knowledge necessary for collective action (Schmid 2005, p. 144).

Although it relies on some unrealistic assumptions, the Coordinated Attack problem does prove that there are instances where full-fledged common knowledge is required for each would-be participant to have rational grounds to contribute to the collective act, and where this is unachievable (Chant and Ernst 2008, p. 552–555).¹⁵ It does not show, however, that common knowledge can never be established by individual agents forming beliefs about each other (nor that communicating intention is never effective in so doing). Using the term “interactive knowledge”¹⁶ for the knowledge one agent has about the knowledge of another

¹⁵ Interestingly, it is the infinite reasoning power ascribed to the two detachments that hinder them: Chant and Ernst show how, for a bounded rationality, at some level of iterated knowledge n , this can no longer be distinguished from $n+1$, and so having reached that level of iteration (however high or low it might be), the two detachments would in fact have achieved common knowledge (Chant and Ernst 2008, p. 564–568).

¹⁶ A term they get from Robert Aumann (Chant and Ernst 2008, p. 550).

(with mutual interactive knowledge constituting common knowledge), Chant and Ernst give an example from the other end of the scale:

Two friends staying in the same hotel in a strange city get separated. It is obvious that they may be perfectly capable of meeting up together again at the hotel rather than, say, at a local landmark. The reason is that this choice is highly salient to both people. For our purposes here, we note that only a minimal level of interactive knowledge is necessary for them to achieve this coordination. Perhaps each needs only believe that the other person is capable of reaching the hotel, for example. No other peculiar features of the situation force them to acquire higher levels of interactive knowledge. (Chant and Ernst 2008, p. 555).

Schmid's use of the Coordinated Attack problem to set the stage for an intentional somersault is, as it stands, partly a straw-man argument. The proponent of an individualistic approach can accept that the two infantry detachments would be unable to fix a time for the upcoming attack in this way, but deny that this is of any great consequence for any but the most artificially constrained examples. In particular, for most if not all real-world situations, there will be other factors influencing belief formation, providing, at least in some cases, rational grounds for contributive action. The real question is not, or rather should not be, whether would-be contributors to a collective act can come to have mutual belief or common knowledge purely through the iteration of beliefs about the intentions or beliefs of the other would-be contributors. As I have argued, the third criterion does not necessitate *any* iteration of beliefs – it suffices that each would-be member of a group forms the belief that the other would-be member(s) intend to contribute. The real question is when or whether individual agents have grounds for forming the required beliefs. I finish by indicating one of the ways the proponent for an individualistic approach might answer Schmid's challenge to explain how individual agents can come to have rational grounds for their contributions to collective acts.

4.2 Grounds for Belief

To recapitulate: At a minimum, the individualistic approach can be seen to posit three criteria for collective acts to occur:

- C1 Each would-be member of a group acts in an appropriately coordinated manner
- C2 Each would-be member intends their behaviour as their contribution to the collective act
- C3 Each would-be member believes that the other would-be members have the intention to contribute to the collective act

Schmid has argued that C3 implies that would-be members must infinitely iterate their beliefs about the intentions of others to contribute, prompting the revised version of C3:

C3* Each would-be member has infinitely iterated beliefs to the effect that the other would-be members have the intention to contribute to the collective act

So far, I have argued 1) that C3 does not imply C3*, and 2) that C3 only gives rise to an infinite iteration of beliefs in the most artificially constrained examples, where iteration also proves to be insufficient for establishing rational grounds for contributing to the collective act. The proponent of an individualistic approach should simply accept that there are such situations, and deny that this is of any great importance. In most cases, there is no need for infinite iterations, nor for psychological mechanisms to compensate for such iterations. However, both the mental parts *and* the actual behaviour are required, and both (or all) parties must fulfil all criteria, *if a collective act is to occur*. If we are looking purely at the would-be contribution of a single agent, this can occur even if no other agent contributes: the only requirement is that the agent in question actually has the required belief and intention, and shows the requisite behaviour. Take Tuomela's example of two people carrying a table (Tuomela 2005, p. 327–328) and focus on the first part of the collective action, the lifting of the table: If I think that you will lift your end of it, and I begin to lift my own, my act would truly be a *contribution* to the would-be collective act of lifting the table – there is no sense in trying to say that I did not do my part. The collective act fails because you did not contribute, not because I was wrong in something. Of course, if this happened, it would quickly become obvious that I was mistaken in thinking that you would do your bit, but that does not alter the fact that my action fulfilled the criteria for my contribution to the collective act. If, *ceteris paribus*, you had contributed, the collective act of lifting the table would have occurred, with zero change in my part of that act. Unless there was some obvious or immediate reason for me to think that you would not contribute, I neither could be sure, nor needed to be sure about your intentions; all that was required was something like a reasonable confidence on my part about you doing yours.

When it is no longer required of each agent that they be *certain* that the other(s) will do or are intending to do their part – or that they be certain that they have no contradictory beliefs about this – the problem of mutual belief is reduced to a question of whether or when agents have grounds for forming such beliefs. But here a whole host of different kinds of evidence can come into play, and we are no longer locked inside the head of the individual contributor (see, e.g. Bratman 2014). The ongoing behaviour of the other (potential) contributors,

any communication of intention or belief that they might make, their responses to your behaviour and communication – even solipsistic mind-games of iterated beliefs, as well as background knowledge or information – can all influence belief formation. While the possibilities are endless, this should not be seen as a weakness of the approach. The simple fact that people tend to be able to act collectively guarantees that there are practicable limits to the kinds of things that go into belief formation in such cases, and this way of approaching the question brings into view the complex sets of factors that *might* influence an individual in an instance of collective action.

This suggestion is also in line with what, e.g. Robert Sugden has already argued about the ways in which individual agents can come to have “reason to believe” that other agents will use “team reasoning” in their interactions (Sugden 2003, p. 176–179), thus making it logically consistent and/or rational for each individual to apply team reasoning. Sugden’s ‘assurance condition’ and to some extent also Michael Bacharach’s ‘interdependence hypothesis’ (2006) can be seen as formalised versions of the kind of belief-forming procedure suggested here. Both have been integrated into theories of collective action demonstrating that common knowledge or mutual belief, while reliant on belief about the beliefs of others, at least does not depend on infinite iterations of belief (e.g. Gold and Sugden 2007 and Tuomela 2013).¹⁷

4.3 Objection: Social Environment

Here it might be objected that we are covertly importing collectivity into the individual analysis, since the grounds for belief stem from something like a pre-existing *social environment*, e.g. in form of the shared norms and knowledge of the individuals involved in collective action.¹⁸ According to this objection, if we think, in Schmid’s example, that the army detachments in most normal cases *would* have rational grounds for launching their collective attack, we do so because we (implicitly) appeal to things such as their belonging to the same army, or the existence of a pre-established chain of command. But these elements may be seen to carry in them the very collectivity or common knowledge that Schmid argues that individual belief-formation presupposes. And so the individualistic analysis ends up depending on collectivity after all.

This objection will not work as it stands, for it presupposes that whatever pre-established collectivity we might appeal to in explaining individual belief

¹⁷ Thanks to the two anonymous reviewers for suggesting this link.

¹⁸ Thanks again to anonymous reviewers for pushing for clarification on this point.

formation in particular cases *itself* necessarily cannot be the result of individual contributions that are not dependent on such collectivity. It trades on a highly constrained and contrived example to draw a general conclusion about all forms of collective action and thus assumes the consequent; that no forms of collective action can be given an individualistic explanation. Whether in the end all forms of collective action presuppose some form of pre-established collectivity remains an open question, but it cannot be assumed against the individualistic approach that this is the case.

5 Conclusion

This article has argued 1) that an individualistic approach to explaining collective action need not give rise to infinite iteration of beliefs, and that 2) freed of the threat of this implausible result, the rationality of individual contributions to collective action can be explained by appeal to a wide range of evidence providing grounds for the requisite beliefs.

The appearance in the individualistic approach of infinite regress, which leads Schmid to conclude that this approach is psychologically unrealistic, results from a conflation of an intention considered as had by Bertha with that intention considered as the object of Anna's belief. Existence conditions pertaining to the first do not transfer to the second, and so one cannot simply assume that the existence conditions for the belief and the intention are shared. One cannot assume, that is, that Anna, in order to form a belief about Bertha's intention to contribute, must also form the same beliefs that Bertha must form in order to have that intention to contribute. With this assumption blocked, the infinite iteration simply does not start. There might be reasons for which one would require of agents that they form beliefs about the beliefs of the other would-be contributors to a collective act, or even that they iterate their beliefs about this, but even this would not occasion any infinite iteration of beliefs.

Collective action does not require of individual contributors that they have indubitable knowledge of the actions or intentions of the other potential contributors, nor that they be certain that they are not acting on a contradictory set of beliefs. We must expect some measure of epistemic responsibility of agents in order to call them rational, but whether someone has rational grounds for contributing to a collective act depends, in all but the most artificially constrained examples, on far more than an iteration of beliefs. A wide variety of evidence can enter into the grounds for forming beliefs, including the existence of a "social environment" of pre-established groups or shared norms and knowledge. This

does not beg the question against Schmid's "Copernican revolution," since the social environment itself might plausibly be constructed from less demanding and less complex interaction between individuals who do not share such an environment.

Whether the individualistic approach to explaining collective action in the end is correct is a question left untouched here. Schmid's suggested alternative opens up a range of new ways to engage the problem of collectivity, pointing in particular to the importance of social environment, and is no doubt deserving of further discussion. But his critique of the individualistic approach does not succeed in showing that this is in need of replacement.

Acknowledgements: Thanks to Leo Townsend, CSMN Oslo, for extensive comments on an earlier draft of this paper. Thanks also to anonymous reviewers and the editors at JSO for their patience and help in bringing this paper to its present state.

Bibliography

- Bacharach, Michael (2006): *Beyond Individual Choice: Teams and Frames in Game Theory*. Princeton, NJ: Princeton University Press.
- Bratman, Michael E. (2014): *Shared Agency: A Planning Theory of Acting Together*. Oxford: Oxford University Press.
- Chant, S. R. and Z. Ernst (2008): "Epistemic Conditions for Collective Action". In: *Mind* 117. No. 467, p. 549–573.
- Gilbert, Margaret (1989): *On Social Facts*. Princeton, NJ: Princeton University Press.
- Gold, Natalie and Robert Sugden (2007): "Collective Intentions and Team Agency". In: *The Journal of Philosophy* 104. No. 3, p. 109–137.
- Heal, Jane (1978): "Common Knowledge." In: *The Philosophical Quarterly* 28. No. 111, p. 116–131.
- Kutz, Christopher (2000): "Acting Together." In: *Philosophy and Phenomenological Research* 61. No. 1, p. 1–31.
- Lewis, David (1969): *Convention: A Philosophical Study*. Cambridge: Harvard University Press.
- Schmid, Hans Bernhard (2005): *Wir-Intentionalität*. Orig.-Ausg. Alber-Reihe Praktische Philosophie. Freiburg; München: Alber.
- Schwitzgebel, Eric (2014): "Belief". In: Edward N. Zalta (Ed.): *The Stanford Encyclopedia of Philosophy*. Spring 2014. <http://plato.stanford.edu/archives/spr2014/entries/belief/>.
- Sugden, Robert (2003): "The Logic of Team Reasoning". In: *Philosophical Explorations* 6. No. 3, p. 165–181.
- Tuomela, Raimo (2005): "We-Intentions Revisited". In: *Philosophical Studies* 125. No. 3, p. 327–369.
- Tuomela, Raimo (2013): *Social Ontology: Collective Intentionality and Group Agents*. New York, NY: Oxford University Press.

- Tuomela, Raimo and Kaarlo Miller (1988): “We-Intentions.” In: *Philosophical Studies* 53. No. 3, p. 367–389.
- Vanderschraaf, Peter and Giacomo Sillari (2014): “Common Knowledge”. In: Edward N. Zalta (Ed.): *The Stanford Encyclopedia of Philosophy*. Spring 2014. <http://plato.stanford.edu/archives/spr2014/entries/common-knowledge/>.
- Weber, Max (1922): *Wirtschaft und Gesellschaft*. Tübingen: Mohr.