Synthese
https://doi.org/10.1007/s11229-019-02355-x

Shared intentions, loose groups, and pooled knowledge

Olivier Roy[1] · **Anne Schwenkenbecher**[2]

Received: 17 April 2018 / Accepted: 31 July 2019

Abstract
We study shared intentions in what we call loose groups. These are groups that lack a codified organizational structure, and where the communication channels between group members are either unreliable or not completely open. We start by formulating two desiderata for shared intentions in such groups. We then argue that no existing account meets these two desiderata, because they assume either too strong or too weak an epistemic condition, that is a condition on what the group members know and believe about what the others intend, know, and believe. We propose an alternative, pooled knowledge, and argue that it allows formulating conditions on shared intentions that meet the two desiderata.

Keywords Shared intentions · Collective action · Pooled knowledge · Common knowledge · Social epistemology

1 Introduction

Many groups that are agents of social change lack a codified organizational structure, and the communication channels between group members are either unreliable or not completely open. Let us call such groups "loose groups". Mobs and social movements are loose groups, as are the anti-globalisation and the "Occupy" protesters (Jordan and Taylor 2004; Bennett et al. 2014), the Arab Spring participants (Allagui and Kuebler 2011) and online activists such as the hackers' group Anonymous (Lievrouw 2011; Olson 2012; Coleman 2014). As these examples suggest, large groups will typically

✉ Anne Schwenkenbecher
a.schwenkenbecher@murdoch.edu.au

Olivier Roy
olivier.roy@uni-bayreuth.de

[1] Department of Philosophy, University of Bayreuth, Universitätsstraße 30, 95447 Bayreuth, Germany

[2] College of Arts, Business, Law and Social Sciences, Murdoch University, 90 South St, Murdoch, WA 6150, Australia

Published online: 08 August 2019

fall under our definition of loose groups, unless they have a codified organizational structure. Small groups may be loose as well when some members can only communicate to others through a third party or through a noisy channel. So, again, we call a group of agents *loose* whenever two conditions obtain: first, the group lacks a codified organizational structure, and second, the communication channels between group members are either unreliable or not completely open.

The paper starts by formulating two desiderata that any satisfactory account of shared intentions in loose groups must meet. These desiderata are based on the observation that, in certain circumstances, shared intentions in loose groups are capable of fulfilling the typical functions of intentions, either individual or shared (Bratman 1999): facilitating interpersonal coordination and planning as well as providing a framework to structure bargaining on shared goals. From that point of view, a satisfactory account of shared intentions in loose groups should be inclusive, and thus weak enough to explain how intentions can be shared in such groups in the first place, but also strong enough to explain how, when possible, they fulfill these functions.

We argue that no existing account meets both our desiderata, and then show, by adapting Bratman's theory of "modest sociality" (2014a, b), how they can be met. By doing so we do not aim at providing a new definition of shared intention. Nor do we want to account for how shared intentions are successfully attributed to groups in natural language. The latter might indeed support cases where the typical functions of intentions are *not* fulfilled, cf. (Kutz 2000; Ludwig 2016), and thus lead to a much weaker set of conditions than those we study in Sect. 7. Rather, again, our aim is to explain how members of loose groups can share intentions in a way that facilitates interpersonal coordination and planning and also provides a framework to structure bargaining on shared goals. We argue that fulfilling these functions has implications for the phenomenology of shared intentionality as well as for the attribution of shared moral responsibility.

Our main point of contention with existing accounts is the epistemic condition on shared intentions (Chant and Ernst 2008). This is the condition regarding what group members know or believe about each other. Many existing accounts use common knowledge as their epistemic condition: among the set of conditions that are necessary or sufficient for sharing an intention is one that states that some or all of the other conditions are common knowledge. As we will see, this is problematic for loose groups. We argue that, instead, what we call *pooled* knowledge allows for existing accounts to be amended so that they adequately apply to loose groups. What will emerge from this discussion is that epistemic conditions may be seen as conducive to, rather than constitutive of, shared intention. Different epistemic conditions support different forms of shared intentions. This, in turn, explains the different roles that shared intentions play in collective actions.

2 Thick and thin accounts of shared intentions

Let us call "thick" an account of shared intentions that makes use of a common knowledge condition, and call the account "thin" otherwise. The accounts of Alonso (2009), Bratman (2014a, b), Cohen and Levesque (1991), Gilbert (1990), Miller (2001), Pater-

notte (2014), Pettit and Schweikard (2006), Shapiro (2014), and Tuomela and Miller (1988) are all thick in this sense. According to these accounts, some or all of the conditions that must otherwise be met for intentions to be shared in a group must also be common knowledge among the group members. Of course, with the exception of Shapiro's all of these accounts can be said to be thick in a further sense: they involve demanding conditions on either the modes of the intentions of the agents, their content, the way the individual intentions interlock with each other, or the personal and interpersonal commitments to the joint activity. In Shapiro's account (2011, 2014), the conditions on the group members' intentions are indeed comparatively weaker. This, however, does not affect our point below, which rests only on the use of a common knowledge condition.

Kutz (2000), Ludwig (2007, 2016), and Pacherie (2013), in contrast, do not feature a common knowledge condition, so we call them "thin accounts". Here again, a terminological clarification is in order. Pacherie (2013) distinguishes "heavyweight" and "lightweight" accounts of shared intentions in terms of the cognitive demands that they make. To the extent that common knowledge is cognitively demanding, our distinction between thick and thin accounts parallels hers. We use a different terminology to emphasize that ours is based only on whether the account involves a common knowledge condition. Thick accounts can be thinned by using a weaker epistemic condition; we will show an example of this in Sect. 7. These thinned versions would presumably still be heavyweight in Pacherie's classification.[1]

3 The common knowledge condition

So, how is common knowledge defined in thick accounts? Here it is useful to distinguish three definitions of common knowledge: classical or iterative, Lewisian, and "ordinary" common knowledge (Paternotte 2017). Let us start with the latter. Ordinary common knowledge is a pre-theoretical concept. It is supposed to capture the fact that something is completely public or "out in the open" in a group. Few thick accounts have moved beyond this. Those that do have mostly used the classical or iterative definition of common knowledge. A proposition P is common knowledge in this sense if and only if everyone in the group knows that P, everyone *knows* that everyone knows that P, and so on, for any iteration of 'everyone knows that'. It is the definition that, for instance, Tuomela and Miller (1988), Schweikard and Schmid (2013), and Bratman (2014a) allude to when they attempt to flesh out the metaphor of P being "out in the open" and move beyond the ordinary, pre-theoretical account of common knowledge. Other definitions of common knowledge exist, for instance, in terms of fixed-point or shared situations (Barwise 1987; Vanderschraaf and Sillari 2014). These, however, imply the iterative definition. So if, as we are going to argue later, common knowledge in the iterative sense is unattainable in loose groups, this applies to these alternative definitions as well.

[1] Note also that for reasons similar to Pacherie's, Butterfill (2012) provides an account of joint *action* through shared goals, which is explicitly aimed at avoiding the use of the notion of shared intentions, either thick or thin.

Lewis' (1969) definition of common knowledge is, in comparison, weaker (Cubitt and Sugden 2003). Instead of entailing the infinite hierarchy of "everyone knows that …", it entails only an infinite hierarchy of "everyone has a reason to believe that …". Classical common knowledge only follows under strong assumptions about the group members' rationality (Lewis 1969). Paternotte's (2014) is, to our knowledge, the only recent thick account of shared intentions that uses Lewisian common knowledge as epistemic condition. We will come back to it below.

4 Shared intentions in loose groups: first desideratum

Classical common knowledge is typically unattainable in loose groups, precisely because communication is either unreliable or not fully open (Halpern and Moses 1990; Parikh and Krasucki 1990). These results do not require that communication is impossible, but only that the group members have no means to make what is called a public announcement in the group (van Benthem 2011). These are announcements to all group members that are completely public, and recursively so: not only does everyone simultaneously receive the message, but everyone also simultaneously receives the information that the message is completely public. In most loose groups listed in the introduction, as well as in our stylized example below (*Coordinated Painting*), such public announcements are not possible.

Therefore any account that uses classical common knowledge as epistemic condition within a set of *sufficient* conditions for shared intentions, as for instance Tuomela and Miller (1988), Schweikard and Schmid (2013), and Bratman (2014a), will not apply to loose groups. The set of conditions that are identified as jointly sufficient for shared intentions will never obtain in such groups. Of course an account that provides only sufficient conditions is not committed to the claim that the conditions put forward, and in particular the common knowledge condition, are individually *necessary* for sharing intentions. But, since common knowledge is not attainable in loose groups, *if* such groups can share intentions *à la* Bratman this will be on the basis of either a subset of those conditions or an altogether different set. It is in this sense that we say that accounts like Bratman's do not apply to loose groups. This does not exclude the possibility of modifying these accounts to arrive at a set of similar conditions that loose groups can meet and that is sufficient for sharing intentions. This is in fact precisely what we do in Sect. 7. Our point, for now, is that loose groups cannot meet all the conditions provided by an account that uses common knowledge as part of a set of sufficient conditions for shared intentions.

If, on the other hand, the classical common knowledge condition is part of a set of *necessary* conditions, as in all the other thick accounts mentioned above, then we are forced to the stronger conclusion that members of loose groups can *never* share intentions. This is problematic, as the following example shows:

> (*Coordinated Painting*) Ann and Bob are painting a house together under the supervision of Claire, the forewoman. Claire has contacted them individually and assigned some tasks to each without the other being present. She has told them (i) that they are part of a team of painters, (ii) that she has told the others

that they are part of a team, and (iii) that she will do the coordination. Ann and Bob have no reason to doubt her word, even though they cannot verify it. Let us furthermore assume that they cannot communicate directly with each other. They go about their individual tasks on the assumption that Claire is indeed taking care of coordinating them, which is what she is in fact doing. She gathers reports of the work's progress from Ann and from Bob and gives each instructions on how to move forward accordingly. The three of them proceed this way and the house gets painted.

We take it to be intuitively uncontroversial that Ann, Bob, and Claire can share the intention that they (*de dicto*[2]) paint the house. This intuition can be supported by observing that, because of Claire's role as coordinator, these intentions can play the functional roles that such mental states, individual or shared, typically play in our mental economy (Bratman 1999): coordinate activities, coordinate planning, and provide a framework to structure bargaining. Observe, however, that if classical common knowledge of each other's intentions is necessary then we are forced to conclude that Ann, Bob, and Claire cannot share the intention that they paint the house. They indeed constitute a loose group in our sense. Ann and Bob can only communicate through Claire. This entails that there can be no classical common knowledge between them. This, again, is a consequence of the results by Halpern and Moses (1990) and Parikh and Krasucki (1990). In 'Appendix 1', we illustrate this argument in detail. For now it is sufficient to observe that, if common knowledge of each others' intentions is necessary, Ann and Bob cannot share intentions either.

Thick accounts using classical common knowledge thus fare poorly when it comes to loose groups: either they do not apply to loose groups, or they yield the intuitively wrong result that it is impossible to share intentions in such groups. It should be noted that this criticism of the common knowledge condition is different from Pacherie's (2013). There, she argues against it on the basis that sharing intentions is possible even for groups of young children, which do not have the (meta-)cognitive resources to attain classical common knowledge. In our example, cognitive limitations are not at stake. We can assume that Ann, Bob, and Claire have the cognitive resources to reach classical common knowledge. The problem lies in the indirect communication channel between Ann and Bob. Blomberg (2016a) also puts forward a plausible example of shared intentions *à la* Bratman, but without classical common knowledge. His case, however, is not one of a loose group in our sense: Hector and Celia can communicate directly with each other. So even though it does not obtain in this case, common knowledge of their intentions is in principle attainable.

Let us call "inclusive" a set of conditions for shared intentions that can be jointly met by loose groups and that yields the intuitively correct answer in cases like *Coordinated Painting*. An account of shared intentions in loose groups should obviously be inclusive in that sense:

[2] By "de dicto" intentions we mean intentions of the form "I intend that we *A*", but where the participants might not know the identity of all the group members included in the "we". They intend that they, whoever "they" are, achieve the joint activity *A*. So *de dicto* intentions stand in contrast with the standard case of "de re" intentions, in which the participants know who the others are.

> **[Desideratum 1]** A satisfactory account of shared intentions in loose groups should provide conditions that loose groups can in some circumstances fulfill, and should yield the intuitively correct answer in cases like *Coordinated Painting*.

Shapiro's (2014) criticism of Bratman's account—essentially that it supposes a level of individual commitment to the joint activity that is never achieved in large-scale collective action—can also be seen as pointing out a failure of inclusiveness. Note, however, that because it itself uses a common knowledge condition, assumedly using classical common knowledge, Shapiro's account also fails in that regard.

Desideratum 1 is however not enough. It requires from prospective accounts of shared intentions in loose groups only a minimal amount of inclusiveness. An overly broad or permissive account would satisfy this without necessarily being satisfactory. This, we argue next, is a problem for the thin accounts of Kutz (2000), Ludwig (2007, 2016), and Pacherie (2013), as well as for Paternotte's (2014) account, which uses Lewisian common knowledge.

5 Shared intentions in loose groups: second desideratum

Let us start with thin accounts. They, of course, meet Desideratum 1, but they yield a too coarse-grained perspective on shared intentions in loose groups. In other words, they blur out important differences in the ways intentions can be shared. To see this, consider the following case:

> (*Carbon Footprint*) Suppose that it is publicly known that if enough individuals change their behaviour relating to greenhouse gas emissions in the right way they will be able to close the so-called emissions gap.[3] No individual is capable, however, to communicate with a sufficient number of others to coordinate this mass action, and each considers it highly unlikely that enough others will spontaneously decide to reduce their carbon emission with the intention of playing their part in the overall plan. Suppose, however, that a large enough number of individuals do so, and the emissions gap is closed.

Ludwig (2016, pp. 220–221) presents a case that has a similar structure. There, two survivors of a nuclear strike manage to launch a retaliation strike despite having no means of communication and each considering it highly unlikely that the other has survived. We use *Carbon Footprint* instead because it yields clearer intuitions in the discussion of moral responsibility below.

On Ludwig's account, these individuals can share the intention to close the emissions gap, even if they do not believe that enough others intend that they fill the gap.[4] There is, *a fortiori,* no common knowledge of these intentions. It suffices that each "intends to contribute to their doing something together in accordance with a common plan"—here, reducing their individual carbon footprint—"where this requires that when acting they all are acting with the same collective action plan in mind" (2016,

[3] Cf. Blok et al. (2012), UNEP (2015), Wynes and Nicholas (2017) for evidence to that effect.

[4] They do not, however, believe that no one else has the intention. This would entail a violation of the belief-consistency condition on intention (Bratman 2009).

p. 290). While it is arguably necessary here that each potential participant knows what Ludwig calls the "common plan", no mutual or higher-order knowledge of what the others intend is required for sharing the intention.

Kutz's (2000) and Pacherie's (2013) accounts lead to similar results. Kutz requires that the individuals have participatory intentions, which are constituted by an individual role (here, reducing their carbon footprint) and a collective end (here, closing the emissions gap). He suggests furthermore a weak epistemic condition, mutual openness, but note that even this might be "overridden" by strategic considerations (2000, p. 19). This condition is essentially stating that the respective participatory intentions would not be undermined if they were out in the open. The subjunctive conditional in this epistemic condition is of course compatible with the group members having in fact no knowledge or belief at all regarding the intentions of others, as is the case here.

Pacherie's account, inspired by Bacharach's (2006) theory of team reasoning, requires instead that, (1) each has "a self-conception as a member of the team" (p. 1833), (2) each believes (1), (3) each identifies what is the best action A for the team, (4) they all, individually, intend to play their part in A. In this case, the epistemic condition (2) is *de dicto*: each individual believes that any others, whoever they are, if any, will conceive of themselves as members of the team that tries to close the emissions gap. There is crucially no belief or knowledge requirement for (3) and (4), but this does not prevent these individuals from sharing the intention to close the gap.

We can grant for our argument that these thin accounts succeed at pinpointing *different*, minimal forms of shared intentions. Indeed the three accounts are not equivalent. Pacherie's conditions, for instance, are compatible with cases which violate Kutz's mutual openness condition. Take, for instance, an intelligence operation which requires complete anonymity, even amongst themselves, of who the operatives are. So already amongst these thin accounts there is not a unique, minimal core of necessary and sufficient conditions for shared intention. Our argument is perfectly compatible with this pluralist view that sees shared intention as a multifaceted phenomenon.

Independently of that, however, it should be clear that each of these accounts is inclusive enough to meet our first desideratum. Loose groups can meet the conditions that these accounts put forward, and they would give the intuitively correct answer that intentions can be shared in cases like *Coordinated Painting*.

Thin accounts, however, do not allow us to distinguish between different ways in which intentions can be shared in loose groups. In particular, they do not help us differentiate between "minimalist" (Kutz 2000, p. 17) shared intentions, or "minimal cooperation" (Paternotte 2014), and fully interlocking shared intentions. The problem is that, in cases like *Carbon Footprint*, the shared intentions can hardly fulfill the typical functions of intentions: coordinate activities, coordinate planning, and provide a framework to structure bargaining. To see this, compare *Carbon Footprint* with a case where the individuals can communicate with each other, possibly through a team of coordinators acting much as Claire does in *Coordinated Painting*. Let us call this variation *Coordinated Carbon Footprint*. Here, the participants' intentions may facilitate interpersonal coordination and planning, and provide a framework to engage, perhaps indirectly, in bargaining. They cannot do so in the original scenario, where participants can only communicate with a very limited number of other potential participants and as such would not be in a position to coordinate their overall actions.

Furthermore, the shared plan is unalterable and there is no place for bargaining. In short, minimal shared intentions cannot fulfill the usual functional role of intentions.

This point is already observed by Butterfill (2012, Sect. II) and Kutz (2000, Sect. V). Kutz claims that this only shows that these roles, and their correlated mutual responsiveness, are not constitutive of shared intentions. As noted earlier, we can grant this for our argument. But, unlike spontaneous crowds or even social movements, which might fall under Kutz's account, there are loose groups where intentions can play these roles. Thin accounts do not explain how this is possible.

The fact that shared intentions cannot play these functional roles has important consequences for the phenomenology of joint intentionality and the attribution of responsibility to group members. First, the joint endeavor will be phenomenologically different in the original *Carbon Footprint* case and its coordinated variant. The participants' sense of joint agency will be weaker in the former than in the latter. Experimental results show that this sense of joint agency is strongly correlated with "not just self-prediction but also … other-predictions and … the joint predictions resulting from the integration of both self- and other-predictions" (Pacherie 2012, p. 373). In minimalist cases such as *Carbon Footprint*, the individuals cannot predict what the others will do, nor what will be the result of their actions. In fact, given that they each think it is highly unlikely that a sufficient number of others will participate, they will experience reducing their individual carbon footprints as a shot in the dark, and its success as a matter of luck, rather than as joint action. The presence of coordinators, in contrast, will provide a better basis for predicting what the others will do and predicting the consequences of joint action, resulting in a stronger sense of joint agency. Our claim here is of course not that the existence of a coordinator is the only factor that is conducive to this stronger sense of shared agency in a case like *Coordinated Carbon Footprint*. Many other specifics of the situation could also do that. We simply claim that the existence of a coordinator is one such factor, because it can raise the degree of confidence in the success of the joint endeavor.

The two variants of *Carbon Footprint* also have different implications for moral responsibility. Suppose, to fix intuition, that instead of succeeding, the group fails to reach its goal of closing the emissions gap, say because the plan was too vague. Who is then to blame? Surely they—collectively—could have done better had they had a better plan. But this was the only one they had, and they were not in a position to deliberate collectively to come up with a better one. As already mentioned, the shared intention to mitigate climate change cannot put pressure on collective deliberation on means, because that type of deliberation is impossible in this case. For this reason it seems difficult to blame the individuals as a group for their collective failure to achieve their shared goal. At best, each of them might bear blame for not having done more individually. But this is not to say that they can be blamed, either individually or collectively, for the failure of the collective endeavor.

In *Coordinated Carbon Footprint*, on the other hand, there is a more solid basis for ascribing responsibility for a collective failure to either the group or the participating individuals, precisely because the shared intentions are able to fulfill their normal functions. Because of the coordinators, the individuals have the possibility to dynamically contribute to the plan as they proceed, possibly by way of interlocking intentions and the establishment of certain forms of interactive knowledge. Furthermore, they

would be able to better foresee and take action to avoid potential problems in executing their plans. Because of this they—collectively—may be considered the authors of the shared plan and they would be the appropriate objects of blame for its flaws. In other words, because of the fact that in *Coordinated Carbon Footprint* the shared intentions can play their normal functional roles, it might seem appropriate to ascribe shared (or joint) blame to the authors and enactors of the plan.[5]

The thin accounts of Pacherie, Kutz, and Ludwig are thus not sufficiently differentiating to explain this variance in moral accountability and responsibility. According to each of these accounts, the set of facts which explains how the individuals share intentions is the same in both variants of *Carbon Footprint*. But even if one is willing to grant that there can be shared intentions in both of these cases, there seem to be differences in the *ways* these intentions are shared, differences that lead to different ascriptions of moral responsibility. The thin accounts cannot explain this.

The outcome of this is that thin accounts are too minimalist to explain how shared intentions come to play the functional roles that are typical of these mental states, and this has consequences for the phenomenological and moral character of the joint action carried out with these intentions. Of course that was not what these accounts were designed for. They are compatible with stronger conditions on how the group members' intentions interlock, and on what they know and believe about one another. But the point remains that these differences are important for understanding the functioning of loose groups, and a satisfactory account of shared intentions in loose groups should be able to explain them. This brings us to our second desideratum:

> [**Desideratum 2**] A satisfactory account of shared intentions should be able to explain the difference between minimal cases, where intentions cannot play their typical functional role, and other cases, where they can.

Neither existing thick accounts using classical common knowledge nor thin accounts of shared intentions meet both desiderata. The former fail on Desideratum 1, the latter on Desideratum 2. Our main goal in this paper is to show that a third way is possible. By replacing common knowledge with a weaker epistemic condition, which we call *pooled knowledge*, existing thick accounts may be amended to meet both desiderata.[6] This is what we will argue for in Sect. 7. First, however, we explain what our proposal, pooled knowledge, is.

6 Pooled knowledge

The epistemic condition that we put forward belongs to the family of so-called aggregate (List 2014) or summative (Gilbert 1989) group attitudes, as opposed to shared or common ones. That is, we are concerned with group knowledge that is "an aggregate or

[5] We appreciate that a more detailed argument would be needed to show exactly how stronger epistemic conditions and shared intentions impact on a collectivity's moral responsibility. We cannot provide such an argument here but will rather simply point to authors who argue that with the right kind of group structure, responsibility may arise at the collective level: May (1992), Isaacs (2011), Held (1970).

[6] In 'Appendix 2', we explain why using Lewisian common knowledge in thick accounts does not solve the problem.

summary of the [knowledge] of the individual members of the collective, produced by some aggregation rule or statistical criterion" (List 2014, p. 1603). Aggregate knowledge is thus the result of putting together the knowledge of the group members, as opposed to expressing what everyone knows, and what everyone knows that everyone knows, and so on.

Distributed knowledge (Halpern and Moses 1990) is the canonical example of aggregate knowledge, and our notion of pooled knowledge rests on this. Distributed knowledge captures the *potential* result of pooling the agents' knowledge. Suppose Aaron and Betty are two scientists working on climate change, and Aaron knows that *if* the climate is warming *then* there will be an increase in extreme weather events, while Betty knows that the climate *is* warming. It follows from their individual information *put together* that there will be more extreme weather events—if they were to pool what they know they would come to this conclusion. This knowledge is distributed between Aaron and Betty.

Distributed knowledge of P can obtain in very different types of situations, as long as P follows from the knowledge obtained by putting together what each individual in that group knows. A proposition can be distributed knowledge in a group without being known by any individual agent. This is the case in our climate scientists example, where Aaron and Betty each know one of the premises together leading to the conclusion that is distributed knowledge. On the other hand, distributed knowledge may be concentrated in a single group member. One agent's knowing P, however large the group G, is sufficient to make P distributed knowledge in G. P being shared or common knowledge in G implies that P is distributed knowledge. In these cases, *each* member of G knows that P. Putting their information together can only lead to more information. So P would still follow.

Let us call a proposition P *genuinely distributed* knowledge in a group whenever P is distributed knowledge in G and P would still be distributed knowledge in the group even if, keeping the truth of P constant, no group member knew P. To illustrate this, consider the case of Aaron and Betty augmented with a third person, Claudia. Suppose that Claudia knows what Aaron knows (*if* the climate is warming *then* there will be an increase in extreme weather events) and what Betty knows (the climate *is* warming) and concludes that there will be an increase in extreme weather events. Her knowledge entails that this conclusion is also distributed knowledge among the three of them. But it would remain distributed knowledge in the group even if Claudia were to suspend judgment on extreme weather events as well as on either of the premises that led her to that conclusion, because it is already distributed knowledge between Aaron and Betty. This piece of distributed knowledge is then *genuinely* distributed in our sense.

The definition of genuinely distributed knowledge is intended to provide a test as to whether the different pieces of information leading to the conclusion really are spread in the group. This is certainly the case if no one in the group knows P but P is distributed knowledge. Each group member knows something which, put together, would entail P. Therefore P is also *genuinely* distributed in this case. By contrast, in the first version of our example, it is distributed knowledge between Aaron and Betty *that the climate is warming*, but this knowledge is not genuinely distributed. Only Betty knows this fact, but if she did not know it, no one in the group would.

The definition of genuinely distributed knowledge requires consideration of a counterfactual situation where P remains true. This is meant to exclude cases where the agents would fail to know P in the counterfactual situation because that proposition happens to be false there. Among those counterfactual situations where the relevant agents do not know P, we should only consider those where their epistemic situation is otherwise minimally changed. One way to capture this idea is to look at cases where the agents have contracted P from their beliefs (see, e.g., Gärdenfors 2003). In such cases, they will typically suspend judgment on P rather than disbelieve the proposition. This of course also entails that in the relevant counterfactual situation the agents might have contracted more than just the belief in P, for instance, if they now know something that entails P. Beliefs whose contents are logically weaker than P, however, should remain. They might in fact positively contribute to keeping the distributed knowledge of P in that counterfactual situation, as in the case of Claudia we have just considered.

The notion of group knowledge that we put forward as an alternative to the common knowledge condition—*pooled knowledge*—is a combination of individual and genuinely distributed knowledge. Let us call a proposition P *pooled knowledge* in group G whenever:

1. at least one member of G knows P, and
2. P is genuinely distributed knowledge in G.

Pooled knowledge is thus a special case of genuinely distributed knowledge where some group members, perhaps only one, know the conclusion P that would result from putting together what is otherwise distributed knowledge in the group. We call these agents the *witnesses* of the pooled knowledge. Of course, if P is known by one or more members of the group then P is distributed knowledge. But as we have seen above, this does not entail that P is genuinely distributed.

The witnesses of pooled knowledge might have come to this knowledge as a result of actively gathering knowledge distributed in the group, possibly putting together the different pieces of the puzzle held by the other group members, so to speak. Pooled knowledge is a definition of the epistemic *state* of the group, however, rather than a definition of the *process* that leads to this state. We will assume nothing about that process, and in particular we will not assume that the witnesses have acquired their pieces of pooled knowledge by actively inquiring amongst the other group members.

In general, common knowledge is consistent with, but does not entail, pooled knowledge. The latter depends on whether P would still be distributed knowledge in the counterfactual situation where *no* group member knows P, whereas in the case of common knowledge all group members know P. Here is a typical case where this would fail: Suppose that Ann and Bob are listening to the radio together and thereby learn of P. Suppose furthermore that this radio announcement makes P common knowledge between them, and that this proposition is logically independent of every other contingent proposition that each knows. Finally suppose that they could not have learned of P otherwise. Then, in the counterfactual situation where neither of them knows P, this proposition is not distributed knowledge. Consequently it is common knowledge but not pooled knowledge.

Classical common knowledge used in thick accounts of shared intention is however a special, symmetrical case of pooled knowledge. The proposition that is common knowledge in these accounts is a conjunction of propositions about the intentions of each of the group members. To the extent that each group member knows what she herself intends, this knowledge would remain in the relevant counterfactual situation where no group member knows what everyone else intends. This is true of each agent in the group. What each group member intends would thus still be distributed knowledge in the counterfactual situation. In this case, common knowledge is a special case of pooled knowledge.[7]

7 Thinning thick accounts

We now show that our two desiderata can be met by replacing common with pooled knowledge in a canonical, thick account of shared intentions: Bratman's (2014a) "modest sociality" account. The core of the argument is that pooled knowledge is strong enough to put rational pressure on its witnesses to step forward and coordinate the joint activity. Of course this argument does not rule out the possibility that Desiderata 1 and 2 could be met by using a different epistemic condition than pooled knowledge. Our claim is not that pooled knowledge is the *only* means to provide a satisfactory account of shared intentions in loose groups. What we defend is a possibility claim.

At the core of Bratman's account is a set of interlocking, individual intentions regarding the joint endeavor and, being a thick account, it uses common knowledge of these intentions. Bratman's conditions on intentions are as follows:

i. Each group member intends the joint activity.
ii. Each group member intends the joint activity by way of the plural intentions of the others.[8]
iii. Each group member intends the joint activity by way of the meshing sub-plans of the intentions of the others.

In a first move, we can substitute the common knowledge condition with one for pooled knowledge. More is needed, however, namely that the witnesses of pooled knowledge are aware of their privileged position in the group:

iv. Conditions i, ii, and iii are pooled knowledge in the group, and the witnesses of pooled knowledge are aware of this.

This new epistemic condition thus makes two substantial assumptions: the existence of witnesses of pooled knowledge of (i)–(iii),[9] and the requirement that they are aware of their epistemic status within the group. Indeed, genuine distributed knowledge of

[7] It would be worthwhile to look at the relationship between pooled knowledge and distributed capacities in more detail, especially in the context of moral responsibility.

[8] Of course, for loose groups all the references to "the others" in these conditions have to be *de dicto*.

[9] In order for the relevant intentions to be supported, the witnesses of pooled knowledge need not know *exactly* what these intentions are. It is arguably enough that the witnesses know that these intentions are oriented towards a single end, cf. [Blomberg, 2016b]. Condition (iv) is thus strictly speaking too strong. It is nonetheless illustrative of how pooled knowledge can support shared intentions, so we leave it at that here.

(i)–(iii) follows if we assume two things: first, that each individual intends the joint activity by way of the intentions and meshing sub-plans of the others, *de dicto*, and second, that each individual knows what she herself intends. From this, we get that each has knowledge of her individual intentions and belief as described in (i)–(iii). So, each agent has a "piece" of the distributed knowledge that conditions (i) and (iii) obtain, namely the one about herself. Now conditions (i)–(iii) are simply conjunctions of these pieces of information. This is what is known by the witnesses of pooled knowledge whom (iv) requires. This means that even if they did not have that knowledge, i.e., if they did not know that *all* intend as in (i)–(iii), this fact would still be distributed, and so it is genuinely distributed. As such, we propose a set of sufficient conditions for shared intention that is less demanding than Bratman's, while not changing the interlocking character of those intentions.

This thinned-down version of modest sociality meets our two desiderata. Let us first look at Desideratum 1. Bratman's original sufficient conditions include a classical common knowledge requirement, which is unattainable in loose groups. Pooled knowledge, in contrast, can obtain in such groups despite the fact that communication between group members is either unreliable or not fully open. An instance of this are the intentions of the three painters in *Coordinated Painting*, which are pooled knowledge as witnessed by Claire.

Furthermore, the group members can rationally intend as in (i)–(iii) even in the absence of common knowledge of what the others intend (or of who and how many they are) if these intentions are pooled knowledge. The potential problem here is that, absent common knowledge, agents might not have sufficient grounds to believe in the possibility of achieving these intentions, and in particular the intentions in (ii) and (iii) that bear on what the others intend. Note that what is required here is of course not the belief that the joint activity will succeed. It is rather that each group member considers it possible that, first, the joint activity is successful, and that it is so by way of the intentions and meshing sub-plans of the group members. This corresponds to Bratman's (2003) 'belief-consistency' requirement, applied to the intentions in (i), (ii), and (iii), respectively, and not to what he calls the "strong belief condition". The interplay between pooled knowledge and the rational pressure put on the witnesses allows circumventing this problem.

The crucial observation here is that the witnesses are themselves group members, and that they thus have intentions as in (i)–(iii). This, we argue now, entails that they will play a coordinating role in the joint activity. Indeed, the more asymmetrically distributed the knowledge about the members' intentions and sub-plans within the group, the more pivotal the witnesses' the role will be. In some cases, the group members will be able to intend by way of the intentions and meshing sub-plans of the others only if these witnesses become coordinators. Because of their privileged access to knowledge, the witnesses' playing a coordinating role becomes a necessary means to the joint activity. Condition (iv), however, also implies that these witnesses are themselves members of the group. So they, too, must have intentions as described by (i)–(iii). Being aware of their privileged position in the group, they will then realize that they have to take a coordinating role in the joint activity, and if they are rational they will form the corresponding intentions. The awareness condition is of course crucial. Otherwise an arbitrary witness of pooled knowledge might mistakenly believe that

she does not need to coordinate because others will do so, even if she intends as in (i)–(iii).

In order to effectively coordinate the group's activities, the witnesses of pooled knowledge will furthermore have to make their status as coordinators somewhat transparent to the other group members. In *Coordinated Painting*, Ann and Bob would not each have formed the intention that they paint the house had Claire not informed them that she is coordinating their activity. So Claire's witness status, but importantly *not necessarily* their individual intentions, is *pairwise* common knowledge between Claire and Ann, and between Claire and Bob, respectively. This provides Ann and Bob with the evidence required for believing that the joint activity is not only achievable but achievable by way of all the members' intentions and meshing sub-plans. In other words, the interplay between pooled knowledge and the individual intentions of the witnesses provides the required evidence for the group members to rationally intend as in (i)–(iii). Sharing intentions is thus possible in *Coordinated Painting* according to the thinned version of modest sociality, which shows that it meets Desideratum 1.

This transparency condition need not be instantiated as pairwise (classical) common knowledge. Sometimes the means of communication between the group members will make this impossible to achieve. This is the case of many online activism campaigns, for instance, the collective Anonymous (Coleman 2014). Following the arguments in Lederman (2017), one might even be skeptical of the possibility of achieving common knowledge even if the relevant facts are completely public. For these reasons we use the purposely vague formulation "somewhat transparent". It will be up to the witnesses of pooled knowledge, if they are rational, to evaluate how transparent they should make their role in the group and to take what they believe to be the necessary means to achieve this. In some situations, something as weak as Lewisian common knowledge might do, in others, shared first-order knowledge, and in some cases, like *Coordinated Painting*, pairwise classical common knowledge. What is crucial is that pairwise transparency *does not need to be a separate condition* for our thinned version of sociality. It follows from the interplay between individual intentions and pooled knowledge.

The argument to the effect that modest sociality, thinned with pooled knowledge, also meets Desideratum 2 rests on the epistemic conditions that distinguish cases like *Carbon Footprint* and its coordinated variant. Modest sociality with pooled rather than common knowledge remains silent on, but does not exclude, minimalistic cases such as *Carbon Footprint*, because the pooled knowledge condition fails in these cases. The coordinated variant is, however captured by the account. As before, pooled knowledge explains how the intentions put rational pressure on the coordinators. It does so in a slightly different way, though, because there might be more than one coordinator. Indeed, while the individual witnesses of pooled knowledge are under rational pressure to make their role somewhat transparent to the other group members, this group of "others" might comprise coordinators like themselves. These do not need to be coordinated, but some level of coordination *between* the coordinators will be necessary. They will have to be mutually responsive and possibly negotiate various roles with one another, and be willing to adapt and discuss changes of plans as they go along. They will, furthermore, have to communicate that fact to the individuals whom they respectively coordinate, allowing the latter to indirectly contribute to the common plan, and initiate further deliberation on means if necessary. This will result

in a stronger sense in which participants share their agency, which in turn makes them potentially accountable for collective failure. This variant of modest sociality thus meets Desideratum 2 as well.

It is worth noting that the weaker variant we propose here is a genuine generalization of modest sociality. This is so because, in this case, common knowledge is a special case of pooled knowledge. If common knowledge of (i)–(iii) obtains, as it does in modest sociality, each group member is a witness of pooled knowledge. In that case, the rational pressure to coordinate, which otherwise applies to witnesses, translates into pressure on each group member for mutual responsiveness and the disposition to coordinate their respective plans and actions as they go along.

8 Concluding remarks

Thin accounts such as Pacherie's, Ludwig's, and Kutz's show that it is possible to share intentions in a minimal sense, where individuals intend to contribute to a shared (or public) plan but otherwise know little or nothing about other participating agents' intentions. If this is correct, then only very minimal epistemic conditions, if any, are necessary for sharing intentions.

What, then, is the proper role of stronger epistemic conditions for the phenomenon of shared intentions? One idea is that they might not be constitutive of but rather conducive to different forms of shared intentions. If the agents know more about the intentions and beliefs of others, these intentions will be more tightly connected, interdependent, or interlocking. This, in turn, explains how these shared attitudes may more robustly play the functional roles which intentions usually play in our mental economy and how this may ground a stronger sense of shared agency as well as more robust types of moral authorship and of moral responsibility. Existing thin approaches lack the explanatory power to account for this. While they provide a minimal account of shared intentions and agency, they miss these important nuances.

However, it is not necessary to go all the way to common knowledge in order to ground these stronger forms of shared intentions. We have argued that pooled knowledge may in some cases be just as conducive to the more tightly connected shared intentions. In other words, while the epistemic requirements of existing thick accounts of shared intention are too strong, there is at least one plausible alternative, viz. pooled knowledge, which provides the right balance of inclusiveness and explanatory power when it comes to showing how shared intentions in a strong sense—with all their normative implications—can exist in large and dispersed groups. Of course, again, this does not rule out that there might be other plausible alternatives to common knowledge, and it would be worth exploring in future projects which other forms of distributed information can secure the role of interlocking intentions in collective action.

Acknowledgements Thanks to Albert Anglberger, Olle Blomberg, Michael Bratman, Justin Bruner, Hein Duijf, Marcel Heidemann, Dominik Klein, Markus Knauff, Alessandra Marra, Luke Roelofs, Frederik Van de Putte, and all participants of the "Joint Duties" Workshop held in Bayreuth in June 2015 for helpful comments and suggestions on earlier drafts of this paper. Further thanks go to audiences at *Collective Intentionality X* in The Hague (2016), the Australian National University (2015) and the University of

Lisbon (2018) where earlier versions of this paper were presented. Part of both authors' work on this paper has been supported by the DFG program for starting of an international collaboration (RO-4548/3-1). Part of Olivier Roy's work on this project has also been supported by the DFG-GACR joint project "From Shared Evidence to Group Attitudes" (RO-4548/6-1).

Appendix 1: No common knowledge of the relevant intentions in Coordinated Painting

In this appendix we explain in more detail why, even though Ann, Bob, and Claire might have shared intentions in our example above, classical common knowledge of this will fail. There are two ways in which the iterative definition of common knowledge can fail: (i) failure of what we call level-1 shared knowledge, i.e., where some participants lack the required knowledge that the others have the corresponding intentions; (ii) failure of second- or higher-level shared knowledge about these intentions. As regards (i), observe that even though for Ann and Bob the identity (and number) of the other(s) is unknown, this does not mean that first-level knowledge will fail. Claire told Ann and Bob that they are part of a team and that she has told the *others* that they are part of a team. Knowing this, it would not be irrational for each of them to form the (*de dicto*) belief that the others, whoever they are, will also plural-intend *de dicto* that they paint the house. For the sake of the argument, we can assume that this belief meets any further requirement to make it qualify as knowledge.

The failure of common knowledge will occur at some higher level of shared knowledge. We have assumed that Claire told Ann that she has told the other(s) that they are part of a team. However, she did not tell Ann that she told the other(s) that she also told *her* (Ann) that she is part of a team. Nor did she tell either of them that they have all received the *same* message. In this situation it is already debatable whether Ann will have the required belief to support the second level of shared knowledge. For all Ann knows, the others, whoever they are, know that they are members of a team. Reflecting on what Claire told her, however, Ann might come to the conclusion that they might be unsure whether she herself (Ann) knows that she is part of a team. In that case, it is not irrational for her *not* to form the belief that the others know that she knows. So Ann knows that each team member knows that they are part of a team, but she does not necessarily know whether the other team members know that *she* knows. This would, of course, be sufficient to break common knowledge of the intentions to paint the house.

But perhaps Ann does not conclude that the others may be unsure whether she herself (Ann) knows that she is part of a team. Maybe she trusts the others to assume that Claire told Ann that she told them. This is still not sufficient to establish common knowledge of the relevant intentions. To ground shared knowledge at level three, for instance, Ann would have to assume that Claire told the others that she told the others that she told the others that they are part of a team. As the levels of 'Claire told the others that …' increase, Ann and Bob will arguably become increasingly uncertain whether the others will make that assumption. But then, even if the increase in uncertainty is very small at each step, at some point it is bound to prevent either Ann or Bob from

holding the belief required by some higher level of shared knowledge, undermining classical common knowledge.

Appendix 2: Thick accounts using Lewisian common knowledge

As mentioned earlier, classical common knowledge is not the only option that thick accounts can resort to. Although most alternatives in the literature are logically *stronger* than classical common knowledge [cf. again Barwise (1987)], Lewis' is, in contrast, substantially weaker. This is the notion that Paternotte (2014) uses to provide an account of shared intention and collective action that is otherwise close to classical thick accounts but should also apply to mass movements like demonstrations or flash mobs, or to cases of cooperation in one-shot Prisoner's Dilemmas, where communication and reputation effects are ruled out. We will now argue that this account, although being substantially more flexible than classical thick ones, faces the same shortcoming as thin accounts. Although it easily meets our first desideratum, Lewisian common knowledge makes it too weak to meet the second.

First a few more details about Lewisian common knowledge. Its key feature is that it requires the existence of a so-called reflexive common indicator for the proposition P that is commonly known (Cubitt and Sugden 2003). This is essentially an event that "indicates" not only P but also that it indicates P for everyone involved. As it turns out, reflexive common indicators can be weakened to distributed reflexive common indicators—cf. again (Cubitt and Sugden 2003)—, allowing for cases where different group members receive different messages. What matters here is that the existence of such a reflexive common indicator for P does not entail classical common knowledge that P. It "only" entails the infinite hierarchy of "everyone has a reason to believe that … everyone has a reason to believe that P". The more rational the group members, Lewis (1969) claims, the more levels of "everyone knows/believes that …" will follow. Classical common knowledge follows from Lewisian common knowledge only for fully idealized, unboundedly rational agents. More importantly, for now, Lewisian common knowledge is also compatible with *no* higher-level knowledge obtaining at all. For each such level, we need to make additional rationality assumptions.

Lewisian common knowledge is one of three cornerstones of Paternotte's (2011) account of weak collective intentions. The account rests on, first, the existence of a shared goal that is, second, supported by some form of game-theoretic equilibrium under, third, Lewisian common knowledge.

Bracketing the details of this proposal, if we grant with Paternotte that there can be weakly collective intentions in his sense in massive demonstrations, flash mobs, and one-shot Prisoners' Dilemmas, and in particular that Lewisian common knowledge of the relevant facts can obtain in such cases, then his account indeed meets our first desideratum. All those are loose groups in our sense. So the account allows for the possibility that members of loose groups share intentions, at least in a minimal sense. Again granting that Claire's respective announcements to Ann and Bob constitute a (distributed) common indicator, it seems that his account would also give the intuitively correct answer in the *Coordinated Painting* example.

The problem lies with our second desideratum. On the one hand, massive demonstrations or flash mobs, or more generally degenerate cases where Lewisian common knowledge of the relevant facts obtains without *any* higher-level knowledge of the intentions and beliefs of others, share important features with *Carbon Footprint*. *If* there is anything like a shared goal or intention in those cases, it will typically not provide a framework to structure bargaining, thus preventing that mental state to play its typical functional role. As we have argued earlier for thin accounts, this is consequential for the phenomenology of shared agency and for the ascription of moral responsibility. Participants in a spontaneous and transient demonstration might not experience the joint action in the same way as those of a highly coordinated group, and it might be difficult to ascribe responsibility to the participants in such an action, or in a flash mob for that matter, beyond that for their individual contribution.

So we end up in a very similar place as for thin accounts. Paternotte's thick account is inclusive, but too much so for our present purposes. It cannot explain the difference between minimal cases, where intentions cannot play their typical functional role, and other cases, where they can. So it does not provide the kind of compromise between inclusiveness and explanatory power that we are looking for. We still need an alternative.

References

Allagui, I., & Kuebler, J. (2011). The Arab Spring and the role of ICTs introduction. *International Journal of Communication, 5*, 8.

Alonso, F. M. (2009). Shared intention, reliance, and interpersonal obligations. *Ethics, 119*(3), 444–475.

Bacharach, M. (2006). *Beyond individual choice: Teams and frames in game theory*. Princeton: Princeton University Press.

Barwise, J. (1987). Three views of common knowledge. In *Proceedings of the 2nd conference on Theoretical aspects of reasoning about knowledge*, Morgan Kaufmann Publishers Inc.

Bennett, W. L., Segerberg, A., & Walker, S. (2014). Organization in the crowd: Peer production in large-scale networked protests. *Information, Communication & Society, 17*(2), 232–260.

Blok, K., Hohne, N., van der Leun, K., & Harrison, N. (2012). Bridging the greenhouse-gas emissions gap. *Nature Climate Change, 2*(7), 471–474.

Blomberg, O. (2016a). Common knowledge and reductionism about shared agency. *Australasian Journal of Philosophy, 94*(2), 315–326.

Blomberg, O. (2016b). Shared intentions and the doxastic single end condition. *Philosophical Studies, 173,* 351–372.

Bratman, M. (1999). *Faces of intention: Selected essays on intention and agency*. Cambridge: Cambridge University Press.

Bratman, M. (2003). Intention, belief, practical, theoretical. In S. Robertson (Ed.), *Spheres of reason: New essays in the philosophy of normativity*. Oxford: Oxford University Press.

Bratman, M. (2009). Intention, belief, and instrumental rationality. In D. Sobel & S. Wall (Eds.), *Reasons for action*. Cambridge: Cambridge University Press.

Bratman, M. (2014a). *Shared agency: A planning theory of acting together*. Oxford: Oxford University Press.

Bratman, M. (2014b). Rational and social agency: Reflections and replies. In M. Vargas & G. Jaffe (Eds.), *Rational and social agency: Essays on the philosophy of Michael Bratman*. Oxford: Oxford University Press.

Butterfill, S. (2012). Joint action and development. *Philosophical Quarterly, 62*(246), 23–47.

Chant, S. R., & Ernst, Z. (2008). Epistemic conditions for collective action. *Mind, 117*(467), 549–573.

Cohen, P. R., & Levesque, H. J. (1991). Teamwork. *Noûs, 25*(4), 487–512.

Coleman, G. (2014). *Hacker, hoaxer, whistleblower, spy: The many faces of Anonymous*. Brooklyn: Verso Books.
Cubitt, R. P., & Sugden, R. (2003). Common knowledge, salience, and convention: A reconstruction of David Lewis' game theory. *Economics and Philosophy, 19*(2), 175–210.
Gärdenfors, P. (2003). *Belief revision, Vol. 29 of Cambridge Tracts in Theoretical Computer Science*. Cambridge: Cambridge University Press.
Gilbert, M. (1989). *On social facts*. Abingdon: Routledge.
Gilbert, M. (1990). Walking together: A paradigmatic social phenomenon. *Midwest Studies in Philosophy, 15,* 1–14.
Halpern, J. Y., & Moses, Y. (1990). Knowledge and common knowledge in a distributed environment. *Journal of the ACM (JACM), 37*(3), 549–587.
Held, V. (1970). Can a random collection of individuals be morally responsible? *Journal of Philosophy, 67,* 471–481.
Isaacs, T. (2011). *Responsibility in collective contexts*. Oxford: Oxford University Press.
Jordan, T., & Taylor, P. (2004). *Hacktivism and cyberwars: Rebels with a cause?* London: Routledge.
Kutz, C. (2000). Acting together. *Philosophy and Phenomenological Research, 61*(1), 131.
Lederman, H. (2017). Uncommon knowledge, *Mind*, online first.
Lewis, D. (1969). *Convention*. Cambridge: Harvard University Press.
Lievrouw, L. (2011). *Alternative and activist new media*. Cambridge: Polity.
List, C. (2014). Three kinds of collective attitudes. *Erkenntnis, 79*(9), 1601–1622.
Ludwig, K. (2007). Collective intentional behavior from the standpoint of semantics. *Noûs, 41*(3), 355–393.
Ludwig, K. (2016). *From individual to plural agency* (Vol. 1). Oxford: Oxford University Press.
May, L. (1992). *Sharing responsibility*. Chicago: Chicago University Press.
Miller, S. (2001). *Social action: A teleological account*. Cambridge: Cambridge University Press.
Olson, P. (2012). *We are Anonymous: Inside the hacker world of Lulzsec, Anonymous, and the global cyber insurgency*. Boston: Little Brown & Company.
Pacherie, E. (2012). The phenomenology of joint action: Self-agency vs joint-agency. In A. Seemann (Ed.), *Joint attention: New developments* (pp. 343–389). Cambridge: MIT Press.
Pacherie, E. (2013). Intentional joint agency: Shared intention lite. *Synthese, 190*(10), 1817–1839.
Parikh, R., & Krasucki, P. (1990). Communication, consensus and knowledge. *Journal of Economic Theory, 52*(1), 178–189.
Paternotte, C. (2011). Being realistic about common knowledge: A Lewisian approach. *Synthese, 183*(2), 249–276.
Paternotte, C. (2014). Minimal cooperation. *Philosophy of the Social Sciences, 44*(1), 45–73.
Paternotte, C. (2017). The fragility of common knowledge. *Erkenntnis, 82*(3), 451–472.
Pettit, P., & Schweikard, D. (2006). Joint actions and group agents. *Philosophy of the Social Sciences, 36*(1), 18–39.
Schweikard, D. P., & Schmid, H. B. (2013). Collective intentionality. In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* (Summer 2013 Edition).
Shapiro, S. (2011). *Legality*. Cambridge: Harvard University Press.
Shapiro, S. (2014). Massively shared agency. In M. Vargas & G. Jaffe (Eds.), *Rational and social agency: Essays on the philosophy of Michael Bratman* (pp. 257–293). Oxford: Oxford University Press.
Tuomela, R., & Miller, K. (1988). We-intentions. *Philosophical Studies, 53,* 367–389.
UNEP (2015) *Sixth emissions gap report*.
van Benthem, J. (2011). *Logical dynamics of information and interaction*. Cambridge: Cambridge University Press.
Vanderschraaf, P., & Sillari, G. (2014). Common knowledge. In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy* (Winter 2014 Edition).
Wynes, S., & Nicholas, K. A. (2017). The climate mitigation gap: Education and government recommendations miss the most effective individual actions. *Environmental Research Letters, 12*(7), 074024.