

Article

Open Access

Kirk Ludwig

Is Distributed Cognition Group Level Cognition?

DOI 10.1515/jso-2015-0001

Abstract: This paper shows that recent arguments from group problem solving and task performance to emergent group level cognition that rest on the social parity and related principles are invalid or question begging. The paper shows that standard attributions of problem solving or task performance to groups require only multiple agents of the outcome, not a group agent over and above its members, whether or not any individual member of the group could have accomplished the task independently.

Keywords: Distributed cognition; Group cognition; Problem solving; Transactive memory systems; Joint action.

It appears to me that in ... philosophical studies, the difficulties and disagreements, of which its history is full, are mainly due to a very simple cause: namely to attempt to answer questions, without first discovering precisely what question it is which you desire to answer.

–G.E. Moore (Preface, *Principia Ethica*)

1 Introduction

Do groups of problem solvers working together give rise to emergent group level cognition? In a 2010 paper in *Cognitive Systems Research*, “Recognizing Group Cognition,” Theiner et al. (2010, p. 378) claim that “recent studies of group problem solving and group memory ... reveal that specific cognitive capacities that are commonly ascribed to individuals are also aptly ascribed at the level of groups.” The thesis has a long history, and has been gaining renewed momentum in recent years with support from a variety of considerations (French 1979; Wegner 1986; Hutchins 1995; List 2003; Heylighen et al. 2004; Tollefsen 2006; Gureckis

Kirk Ludwig, Indiana University – Philosophy, 1033 E. Third St., Sycamore Hall 026, Bloomington, IN 47401-7005, USA, e-mail: ludwig@indiana.edu

 ©2015, Kirk Ludwig, published by De Gruyter.

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 3.0 License.

and Goldstone 2006; Goldstone et al. 2008; Theiner 2009, 2013a,b; Sutton et al. 2010; Theiner and O'Connor 2010; List and Pettit 2011; Gallagher 2013; Michaelian and Sutton 2013; Szanto 2014).

My goal in this paper is to evaluate some recent arguments with a somewhat novel strategy for the claim that at least certain extant forms of distributed cognition involving groups of cognizers solving problems or performing tasks – that the individuals involved could not do alone or with the same efficiency or facility – are genuine instances of group level cognitive processes, where this implies in particular that *the group itself* can be said to have in some nontrivial sense cognitive capacities, cognitive properties, or cognitive states. I will argue that the considerations advanced in these arguments provide no reason to think that extant forms of distributed cognition involving groups of cognizers solving problems or performing tasks are instances of group level cognitive processes, in any but a *Pickwickian* sense.

We can start by clarifying the question.

What do I have in mind by *distributed cognition*? I have in mind a group of agents who are jointly intentionally engaged in problem solving or task performance, where solving the problem or performing the task requires drawing on their cognitive resources, e.g., their knowledge and skills and abilities to solve problems, and where they each make a contribution to solving the problem or carrying out the task. Distributive cognition is not by definition group level cognition (this mistake is made in (Button 2008), among others, who, though arguing against group minds, does not distinguish this from distributed cognition).

What do I have in mind by *group level cognitive processes*? A cognitive process, in the sense at issue, if the question is to be of any interest, *must involve a transition in cognitive states that have original and not derived or as-if intentionality, and which are an expression of a capacity to bring to bear knowledge and skills in the solution of problems or performance of a task.*¹ To say that there are group level cognitive processes is to say that such processes involve states of the group itself in addition to its members. For example, a group that solves a problem might be said to *itself* be the agent that solves the problem, or to have itself cognitive capacities, analogous to those of individuals, which are its cognitive capacities and not the cognitive capacities of the individuals, and the group *itself*, in virtue

¹ For readers who may be unfamiliar with these expressions, see the entry on “Intentionality” in (Pfeifer and Sarkar 2006). Intentionality is the property of being directed at or about something. Something – an event, state, object, sentence, etc. – has derived intentionality if and only if (iff) its intentionality is explained in terms of the intentionality of something else. Something has original intentionality iff its intentionality is not derived. Something has as-if intentionality iff it does not have intentionality but behaves as if it did in one or another way.

of these things, may be said to be *the subject of intentional states, both representational and goal directed*.

It is in this sense that I will deny that the arguments to be considered give us any reason to think that any extant forms of distributive cognition in problem solving and task performance involve group level cognition.

In the following, I begin with a review of some arguments for group level cognition of the sort I want to focus on. I take as my stalking horse Theiner, Allen, and Goldstone's (henceforth TAG) "Recognizing Group Cognition," which I think is an paradigm of the genre (2010, page citations alone are to this article), and which aims to settle the issue without having a grand theory of the mind or cognition. At this stage, I aim to lay out the arguments of the proponents in as a beguiling fashion as possible, hewing relatively closely to their favored ways of putting things. Then, I will consider each sort of case, and I will argue in detail that nothing of any substance is established by careful reflection on them. I cannot hope to review every example or every sort of argument, but I think the examples I review will show something of a general nature about the way these arguments go and provide the materials for debunking other projects along the same lines.

2 Arguments from Distributed Cognition to Group Level Cognitive Processes

... in psychology there are experimental methods and *conceptual confusions* ... The existence of the experimental method makes us think we have the means of solving problems which trouble us; though problem and method pass one another by.

—Wittgenstein (*Philosophical Investigations* §693)

2.1 Types of Group Problem Solving

There are various ways in which groups of cognizers can work together to solve problems (Laughlin 2011). A classic taxonomy is provided in (Steiner 1966).

1. *Additive*: They can do something that is the sum of their individual contributions, e.g., in a task that requires the group to process credit cards applications by each working independently to process credit card applications.
2. *Compensatory*: They can produce an outcome that is a function of their individual solutions, e.g., they can aggregate judgments by a majoritarian rule, or take the average of the estimates of stockbrokers on the future value of a company's stock.

3. *Conjunctive*: They can do something that requires each of them to succeed at the same thing for all of them to succeed, as in the case of a team of mountain climbers, or a crew of rowers, synchronized swimmers or divers, or, in the case of a problem solving task, a mathematics competition in which inter alia a team is presented with a problem which each of its members must solve in order to pass to the next level of competition.
4. *Disjunctive*: They can work independently and as a group choose the best solution out of the individual efforts. For example, a group may enter a design competition and then have members work independently and choose the best design to represent as the group's design.
5. *Complementary*: They can institute a division of labor in which a complex task is subdivided into different parts which are assigned to different agents, and, indeed, the problem solving can be spread across the agents and instruments and structural features of their environment. Most organization problem solving is of this type – football teams, orchestras, research teams, universities, corporations, and so on.

All of these involve distributed cognition, but so far as I know only compensatory and complementary problem solving have been cited in arguments for group level cognitive processes.

For example, some arguments for group level cognitive processes focus on the conditions for a group to be capable of stable rational judgment aggregation (Pettit 2003; List and Pettit 2011). These are compensatory tasks in the taxonomy above. Given certain assumptions about members of the group, namely, that they make judgments independently, and that they are positively biased toward the truth, one can show that it is more likely, and as the number of members of the group goes up, much more likely, that the judgment that the majority makes is correct than it is that any individual's judgment is taken independently, provided that all of them are also fallible (Grofman et al. 1983). Let us call this, in a case in which members of a group decided to use this as a method for arriving at a decision about what proposition they will act on or accept, the group judgment. Then in appropriate conditions the group's ability to arrive at correct group judgments will be significantly better than that of any individual member of the group. The by now well-known discursive dilemma shows that in more complicated decision making contexts, where multiple judgments must be combined, unless the judgment aggregation process is carefully designed, there is no guarantee that the judgment reached will be rational, in the sense that the group may endorse a set of propositions that are not jointly consistent. List and Pettit have suggested that when a group has a properly designed decision procedure in the sense that it is set up so as to avoid inconsistent results, we have an instance

of a genuine rational group agent. But this will not be my main concern in this paper.

My main concern is with arguments from complementary problem solving tasks to the existence of group level cognitive processes.

2.2 The Argument Strategy

The arguments I have in mind involve an overall strategy, a set of considerations about what group cognition would come to, and some methodological principles, which are applied to a set of examples. If the application of the strategy in light the considerations and principles shows in these cases that there is group level cognition, then group level cognition is pervasive, even if it has long gone unnoticed.

The strategy is to eschew the grand project of giving an analysis of the distinction between cognitive processes, capacities, properties, and states in the abstract, and instead to focus on particular cognitive processes, capacities and properties, and, in particular, on problem solving capacities and the use of memory in group task performance. The idea is to avoid becoming entangled in a general demarcation problem and instead to show that there are group level cognitive properties by showing that determinate forms of cognition can be attributed at the group level. This would suffice for showing the determinable is instantiated whether or not we have a general analysis of how to demarcate the cognitive from the non-cognitive.

As TAG put it:

Some mental properties seem clearly projectible to groups – groups solve problems that individuals cannot, for example – but others, like consciousness, seem equally unprojectible. Because of the heterogeneity among different cognitive and mental predicates, we believe that abstract arguments about group minds or extended minds should be replaced by specific discussions tied to particular properties: group memory, group problem solving, etc. Instead of looking for “the mark of the cognitive” (Adams and Aizawa 2008) we should rather be asking whether specific cognitive models that work at the level of individuals also work at the level of groups. (p. 379)

Our underlying assumption here is that the best way to determine whether groups have cognitive properties is to seek out successful characterizations of group behavior that invoke psychological constructs at the level of groups, and to catalog the empirical and methodological credentials of such characterizations. (p. 385)

Before we turn to the set of considerations and the methodological principles that drive the argument, let’s reflect on the nature of the project itself. Notice how TAG

set up their project. They start by saying that some *mental properties* are clearly projectable to groups (whether their example actually shows that is another question – see Section 3.2 below). They do not suggest that they have in mind an esoteric or technical sense for ‘mental properties’. Moreover, the two traditional marks of the mental are intentionality and consciousness. They say that while some mental properties are projectible, others like consciousness are not. The sorts of mental properties that they have in mind seem to be, as this indicates, not theoretical constructions, but the ordinary notions. They appear in particular to have in mind by cognition a process that involves intentionality (states with representational content). They go on to say that they want to replace abstract arguments about group minds or extended minds with discussions about particular mental properties or processes. The debate about group minds and extended minds is typically about whether or not the entities that have mental states, or undergo mental processes, include groups or individuals-together-with-features-of-their-environment [one might consider the exchange between Fodor (Fodor 2009) and Clark (Clark 2009)]. Thus, they are clearly setting up their project in the context of a debate about whether the subject of mental states and processes in the usual sense are just individuals or include groups as well. And they give as examples of the mental properties they have in mind memory and problem solving, both of which, in the ordinary sense, involve states with original intentionality. They go on to say they are not going to look for the mark of the cognitive, contrasting their goal with that of Adams and Aizawa. They do not want to become involved in the demarcation debate. It is notable, in this connection, that Adams and Aizawa criticize proponents of the extended mind hypothesis for not saying what they mean by a cognitive process, and they argue that saying just that cognition is any process that contributes to intelligent behavior is an exercise in relabeling with no scientific payoff (Adams and Aizawa 2008, p. 73, 85). For example, there are many chemical processes in the brain that are necessary for intelligent behavior – are they *ipso facto* cognitive processes? If this were the idea, it would take the wind out of the sails of the argument. We may assume that TAG do not want to be tarred with this brush. TAG say explicitly that they want to show that “group cognition is neither *trivial* nor shrouded in metaphysical mystery” (p. 379; emphasis added). That motivates in part their focusing on particular non-controversial recognizable determinate sorts of mental properties and processes. If we already accept that certain properties and processes are mental, and we show that those are instantiated in groups, then we can sidestep this criticism and the demarcation debate in one stroke. It is thus quite important that TAG not shift from ordinary notions of cognitive states and processes to some sort that is theoretically loaded and controversial, if their strategy of sidestepping this debate is to be successful.

I belabor this point because one, perhaps natural, response to the criticisms that will follow is to try to move the target out of the way by claiming that the arguments have no ambition to show anything about groups having mental properties in any ordinary sense. One might say, e.g., that the talk of ‘cognitive models’ working at the level of individuals also working at the level of groups, in the last sentence in the passage above, signals that they are *not* interested in processes that involve intentionality. But this interpretation does not fit with the way the project is set up. Talk about group *minds* (a group *having a mind*), and *mental* properties being projectible to groups (not just their members), the contrast with looking for the mark of the cognitive, the idea that we can avoid the problems that Adams and Aizawa raise for the extended mind hypothesis by focusing on particular determinate, uncontroversial instances of *mental states and processes*, the explicit goal to avoid *trivializing* the thesis, none of this would make sense if the idea were to focus on some theoretical notion of “cognitive process” which was detached from the traditional marks of the mental, or associated with them as only being a sort of infrastructure for them.² Furthermore, as we will see, the only two “cognitive processes” they consider are problem solving and memory, both of which involve intentional states.

It is possible that the project is *introduced* in a misleading way, and I will come back to the suggestion that it might be construed differently and what the significance of that might be. The default interpretive assumption, however, is that they are engaged in the project that they represent themselves as engaged in at the outset. For the time being, I proceed on the assumption that TAG are interested in showing that at least some *groups* of individuals have mental states and undergo mental processes just as their members can. It is an *important* question whether extant groups have mental states and undergo mental processes in a non-trivial sense, and it is an important question whether the (representative) considerations advanced by TAG support the thesis or not. This will help to define the limits of what such considerations might establish.³

² The first author, Georg Theiner, provides a more detailed account of how he understands group cognition, in the context of discussing Transactive Memory Systems, in (Theiner 2009). He lists seven properties which contribute to the cognitive nature of a system S: (1) “S can adapt its behavior to changing environments”; (2) “S can process information from its environment”; (3) “S can selectively and purposefully attend to its environment”; (4) “S can create internal representations of its environment”; (5) “S can modify its environment through the creation of artifacts”; (6) “S can become aware of itself as a cognitive agent”; (7) “S can have conscious experiences of itself and the world” (p. 335). He discusses 1–5 in connection with Transactive Memory Systems. One might read (1) and (2) behavioristically, but (3)–(5) obviously involve intentionality. Theiner attributes (1)–(5) to Transactive Memory Systems.

³ There is a danger in discussions about the scope of cognition of a terminological slide or drift (perhaps unwittingly) between talk of minds and mental states, reasoning, and agency, and some-

With this behind us, let us turn to the considerations and methodological principles that underlie the argument. The considerations are these:

Desiderata on group cognition (at least of the relevant type).

- R1. Group Cognition is an emergent phenomenon, at least in the sense that it is not the effect of mere aggregativity.⁴
- R2. Group Cognition involves a differentiation of roles (division of cognitive labor) of the members of the group that involve an organizational structure in the group that explains emergent group level cognition.
- R3. Emergent group level cognition identifies a level of cognitive organization that is also an important explanatory kind with respect to the behavior of the system.

It is important to note that R1–R3 are not introduced as sufficient conditions for group level cognition.⁵ In the case of R1, it is clear that not all emergent phenomena are ipso facto cognition, e.g., wave motion in water. With respect to R2, it is clear that division of cognitive roles is not sufficient for group level

thing rather less exciting, aided by the use of cognates of the term ‘cognition’, and especially ‘cognitive system’, ‘cognitive process’, ‘cognitive capacity’ and ‘cognitive property’, which *can* be read in a robust sense, but are vague enough to be read also in a much more attenuated sense. Perhaps everyone could agree that there is a sense in which people working together to solve a problem are a cognitive system and that they engage in cognitive processing which involves all of them, that they have as a group cognitive capacities (they can do things as a group such as solve problems), and that the group has cognitive properties (e.g., being composed of cognitive agents). They are, after all, doing cognitive work together! A proposal of this sort is not suited to raise either eyebrows or interest. We need to keep in view what could be a substantive claim in the area. We should take care, in particular, not to infer from the anodyne claims that groups are in this pedestrian sense cognitive systems, undergo cognitive processes, have cognitive capacities, and have cognitive properties, to the claim that they (the groups per se) cognize (reason, know, understand), hence, are cognizers (thinkers), where we understand the latter claims to be substantive.

4 “A property P of a complex system S is aggregative if and only if (i) P(S) is invariant under the inter-substitution of parts of S, or any other parts taken from a relevantly similar domain; (ii) P(S) remains qualitatively similar (differing only in value) under the addition or subtraction of parts; (iii) P(S) is invariant under the decomposition and re-aggregation of parts; and (iv) there are no cooperative or inhibitory interactions among parts” (p. 379).

5 So far as TAG provide sufficient conditions, they appear to be expressed here (this is M1 below): “If, in confronting some task, a group collectively functions in a process which, were it done in the head, would be accepted as a cognitive process, then that group is performing that cognitive process [(Theiner 2008); see also (Deborah Tollefsen 2006)]. A potential worry about the use of such principles is that they implicitly appeal to shared intuitions of an unspecified community about the constitutive functions of cognitive processes in the head. Since we do not want to rest our argument on presumptive folk intuitions about cognition, we suggest splitting up the two main epistemic roles which social parity, as well as the original parity principle, were meant to perform. Its first, meta-theoretical role is merely to establish a metaphysical ‘default’ assumption which allows us

cognition. If coauthors of a book divide up chapters, that does not show they are performing group level cognition. And, in any case, R2 only specifies a constraint on emergent group level cognition, since its application presupposes independent identification of group level cognition. R3 likewise is stated as a constraint on emergent group level cognition, and it is furthermore clear not every important explanatory kind with respect to the behavior of a system is emergent group level cognition, such as, e.g., whether the system has gills or lungs.

Methodological Principles

- M1. The social parity principle: “If, in confronting some task, a group collectively functions in a process which, were it done in the head, would be accepted⁶ as a cognitive process, then that group is performing that cognitive process” (Theiner 2008; see also Tollefsen 2006).⁷
- M2. We have evidence for group level cognition when the capacities for problem solving that can be ascribed to the group are distinct from those possessed by any individual members of the group.

to ascertain what burden of proof is appropriate for either side of the debate. Its second role is to provide theoretically relevant criteria for individuating cognitive states and processes. A cogent argument for group cognition must jointly satisfy both requirements, but it need not do so by appealing to a single principle” (p. 384). There is, however, no further specification of the second of these components, and so I will focus on the original social parity principle in the discussion.

6 I will assume throughout that the intention is that this be read as “correctly accepted”.

7 More on the social parity principle: “the social parity principle rests on a fairly generic, intuitive understanding of cognition, in which cognitive properties are identified by the causal-explanatory roles which they play in commonsensical psychological explanations. This understanding may well be sufficient when considering whether groups can have folk-psychological properties such as beliefs, desires, or intentions (Gilbert 1989; Tollefsen 2004). However, because we do not want to bite off more than we have space to chew, and because we wish to remain grounded in actual psychological research, we will not have anything to say about the attribution of propositional attitudes or consciousness to groups” (p. 385). They may suggest an ambition that falls short of showing that there are genuine group level cognitive processes in the sense I have said that I am interested in. They go on to say: “we confine our attention to specific cognitive capacities that have traditionally been studied by psychologists at the level of individuals, but have more recently been used to analyze the behavior of groups as a whole. ... [we] seek ... successful characterizations of group behavior that invoke psychological constructs at the level of groups” (loc. cit.). And talk of “psychological constructs,” as if the idea were to consider some theoretical concept, may reinforce this impression. However, they immediately go on to say: “In this paper, we focus on group problem solving ... and on group memory” (p. 385). So we are back to ordinary intentional processes, capacities and states. As I have noted, will return toward the end to the question what sort of “cognitive capacities” we might see groups as having if we do not see them as having genuine intentionality, and what its significance might be.

M2 is never explicitly stated but appears to be implicit in the following passage, and seems likewise to play an implicit role in the argumentation.

... groups have structure, and this structure is important to their behavior, including their ability to adapt to different circumstances. It is this structure that allows us to speak of mechanisms of group cognition. Group cognition is not simply the unstructured aggregation of individual cognition, but the outcome of a division of cognitive labor among cognitive agents. Such division of cognitive labor may be the result of explicit organizational decisions by the individual agents, or (and we believe more commonly) *the result of interactions among the agents that lead to enhanced group capacities* without the express intent of the agents. (p. 379; emphasis added)

2.3 Examples of Group Level Cognition

With this as background, I turn to some examples of group problem solving which have been advanced to show that there is genuine non-aggregative emergent group level cognition.

1. Stigmergic path formation
2. The Collective Coloring Problem
3. Acquired division of labor in groups
4. Group remembering and Transactive Memory Systems

2.3.1 Stigmergic Path Formation

The first example is stigmergic path formation. Stigmergy is a mechanism of indirect coordination by independently acting agents in which a trace left in the environment by action increases the likelihood of its repetition by the same or a different agent, so that actions build on each other and lead to spontaneous emergence of a group level coherent pattern of activity. The term was originally introduced by the French biologist Pierre-Paul Grassé (Grassé 1959) in connection with the behavioral patterns of social insects. An example of a stigmergic mechanism operating in a human group is the creation of pathways across lawns between university buildings by following the traces left by previous walkers.⁸ What path systems arise depend on a number of factors. Agents find it easier to follow a path that has been visibly left by a predecessor. In doing so, they reinforce the trails, making them more visible, and, hence, more highly attrac-

⁸ (Helbing et al. 1997a,b; Goldstone and Roberts 2006; Moussaid et al. 2009).

tive to subsequent walkers. Walkers also want to take the shortest path to their destinations. Walkers' movements are a compromise between taking the shortest path and taking the easiest path. The system of trails that emerges depends on the values attached to taking the shortest path and the easiest path, the relation of visibility to judgments of ease, the decay rate of the visibility of a trace left by a walker, and the amount of traffic. If an optimal solution for a group is a path system that conserves the total amount of pathway, appropriately increasing path visibility or path decay relative to a certain amount of traffic can lead to an optimal solution, though none of the walkers intends it (Figure 1).

Here, as TAG put it:

The group as a whole can be aptly construed solving a problem – technically the problem of approximating Minimal Steiner Trees defined as the set of paths that connects a set of points using the minimal amount of total path length. This is a classic, hard information processing problem ... and yet good approximate solutions can be found without any cognizant human planner or even any person knowing the problem their group is engaged in solving. ... It is exactly the kind of pattern that makes it plausible that extending information processing concepts to systems other than those composed of neurons will prove fertile and inductively powerful. (p. 386)

Thus: there is emergent group level problem solving because a group solves a problem although no member of the group (or anyone else) even knows what problem the group is engaged in solving (M2), and it is a problem which, if solved




Fig. 1: This shows the influence of different parameters on path formation in the Active Walker model. A is the optimum path system to conserve total pathway length, and is achieved by a combination of high decay and visibility. B differs from A in having $1/100^{\text{th}}$ the degree of visibility. C differs from A in having $1/100^{\text{th}}$ the rate of decay and $1/10^{\text{th}}$ the level of visibility. D differs in having $1/10^{\text{th}}$ the level of visibility (Goldstone and Roberts 2006).

by an individual would count as a cognitive process (M1). We can note further (R1) that the capacity to solve the problem is emergent at least in the sense that it is non-aggregative (since it involves interactions that stimulate or inhibit among the members of the group), and that the level of organization relevant to explaining the behavior of the system (R3) has to do not just with properties of the members of the group but also features of their environment (such as decay of paths and degree of visibility of a trace).

2.3.2 The Collective Coloring Problem

In the collective coloring problem, groups of people are arranged in different social networks, and their individual tasks are to choose their colors so that no network neighbor shares the same color. Thus, in contrast to the previous case, they have a shared goal that they represent themselves as pursuing with the others, that is, they intend jointly to solve the problem. Two topologies are shown in Figure 2.

Groups usually found solutions, though in some topologies it was easier to find solutions than others, e.g., cycle is easier than preferential attachment. Parameters include degree of visibility of others in the network, and an interesting result is that limiting visibility to immediate neighbors gives rise to quicker solutions.

The important points are these:

... the group is able to flexibly solve a class of difficult network coloring problems – problems that, *if they were solved by individual humans*, would be interpreted as requiring relatively sophisticated problem solving strategies including planning, back-tracking, and




Fig. 2: Two topologies for the collective coloring problem.

multiple constraint satisfaction (M1?, M2). Second, the groups' problem solving capacity depends on group-level parameters such as network topology, and is not determined in a straightforward manner from individuals' problem solving capacities (R1, R3?). If it were, then we would generally expect that giving individuals more information about the nature of the problem would facilitate the group's problem solving ability. (pp. 386–387; emphasis added)

Thus (the argument goes), we have genuine group level emergent cognitive processes.

2.3.3 Acquired Division of Labor in Groups

In this sort of paradigm, members of a group share a goal, that is, they want jointly to carry out a task. So, as in the last case, they represent their shared goal and undertake to jointly intentionally pursue it. The object is to run a task over multiple rounds so that the group can adapt how they pursue it and to structure the task in such a way that benefits accrue from a division of roles in the group. The particular task was playing a game called "Group Binary Search". One hundred and six participants were divided into 18 internet connected groups and they were told a computer would select a number randomly between 51 and 100. In each 15 second round, each participant enters a number between 0 and 50 and the numbers for each group are summed and compared to the target sum. The participants (in a group) know what numbers others choose. In case 1, the "numeric feedback" condition, they are told how high or low their group was. In case 2, the "directional feedback" condition, participants are only told whether their group's sum was too high or too low. Then the game is repeated. Five games in each condition were played by the 18 groups. Group performance increased as the number of members decreased. With one participant, in the numerical feedback game, at most two rounds are needed (choose 0 the first round and then the number by which one is low the second round); for the directional feedback condition one does a binary search with at most $\log_2(49)$ rounds. Groups improved over repeated rounds and in part because of a differentiation of roles. Some participants in groups altered their numbers only a bit in light of feedback, others more, but over rounds this tendency increased, and because participants know what numbers others choose, this enables them to predict the behavior of others (Figure 3). In larger groups the changes across people increase over rounds. Thus, there is a differentiation into reactive roles and nonreactive roles. The more diverse, the greater the performance increases. "In sum, the group comes to be able to solve a difficult coordination problem because the variability within individuals' behavior [reactive vs. nonreactive] is decreased, and across individuals it is increased" (p. 387).

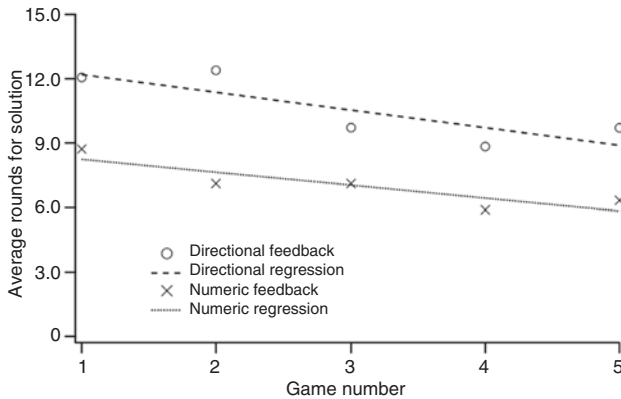


Fig. 3: Improvement in group performance over repeated trials.

What does this show?

... the experiment shows how groups containing individuals that respond to each others' strategies can organize themselves. The potentially socially relevant conclusion to this work is similar to one that (Page 2007) has emphasized – that groups perform better when they are composed from diverse individuals. Page provides proofs that diverse groups will do better in well-delineated circumstances compared to groups that are made up of the best-performing individuals. However, the new perspective from the Group Binary Search experiment is that individuals in a group can learn over rounds and over games to become diverse even if they did not start out that way, and if they can do so, then they can solve group coordination problems more efficiently. (p. 388)

The point appears to be that groups may self-organize in order to solve a problem on which they coordinate more efficiently by differentiating roles. Presumably, the idea is that we can invoke the social parity principle (M1) here, as well as R1, R2, and R3.

2.3.4 Remembering in Groups

The last paradigm is the transactive memory system in which a group performs a task, such as assembling a radio, where there is a division of labor, but also a division in cognitive skills and knowledge (know-how and explicit memory) about how various portions of the task are to be completed. Individuals in transactive memory systems also possess knowledge of who knows what and who has what bit of expertise relevant to the task the group undertakes, and when to call upon

the knowledge and skills of others in performance of the group task. There is evidence that groups that learn their tasks together perform better together than groups formed randomly from those trained for the subtasks. The upshot is supposed to be:

Groups containing members that know each other well, have had a long time to adapt to each other, and trust one another seem to be able to form TMSs that *transcend what individuals are capable of when considered independently*. Moreover, the TMS perspective on memory, cut loose from an individual-centered restriction, helps identify the requirements of any memory system: (A) *information is stored from experiences*, (B) *subsequent access to this stored information*, and (C) *subsequent use of the accessed information to deal with related situations more effectively*. These functional requirements do not require that memories be the exclusive property of individuals, and in fact there are documented cases of collective memories satisfying these requirements better than individual memories. (p. 389; emphasis added)

Thus, again, we have a case for the application of M1, M2, R1, R2, and R3.

In sum, the form of the argument is this: there are cases in which groups collectively function in a process that, if it were done in the head, would be accepted as a cognitive process, where in addition, the group level capacities for problem solving (or task performance) are distinct from those possessed by any individual members of the group, which satisfies plausible (if perhaps defeasible) requirements on group cognition, such as that it is emergent in the sense of not being aggregative, that it involves a differentiation of roles, and that it introduces a level of cognitive organization that is an important explanatory kind with respect to the behavior of the system.

3 Critical Analysis

The trouble is, as so often in philosophy, it is hard to improve intelligibility while retaining the excitement. —Donald Davidson

I turn now to a critical examination of the style of argument just surveyed in application to the examples given.

I begin with some general remarks to frame the issues.

First, the issue before us is not whether in principle a system consisting of multiple cognizers could itself be a cognizer. Whether one is a functionalist or an emergentist, the in-principle possibility of a group of agents that is itself an entity with mental properties is a conceptual possibility. So my intention is not to argue from general principles to the conclusion that no group that engages in problem

solving can be a cognitive agent in its own right. My goal is rather to examine certain arguments that have been advanced to show that particular actual examples of group problem solving and task performance are also instances of group level cognitive processes.

Second, it is clearly relevant to the general issue whether or not functionalism is an adequate theory of cognition, whatever we might think of its promise as a theory of phenomenal consciousness. If functionalism is not an adequate theory of cognition, then the bar for showing that groups have genuine cognitive properties becomes higher. However, it is not my intent to argue that functionalism is not an adequate theory of cognition. The critical points that I make will not depend on that.

Third, I am interested in whether groups which solve problems of which we have extant examples have genuine group level cognitive properties, capacities and so on. I am not interested in whether they act *as if they do*, whether they in fact do or not. Behaviorism is not a viable account of individual cognition, and it does not become one when it is applied to group behavior. Furthermore, we are interested in original intentionality, and not derived or as-if intentionality. We are not interested in whether relative to an assignment by an external observer a system has a state that represents something, but whether it has observer or assignment independent intentionality, and we are not interested in whether something acts merely as if it has intentionality.⁹

Fourth, connected with this, what I have in mind by group level cognitive capacities, processes, properties and so on is group level instantiation of states that have intentionality. We are not interested in states that carry information in the sense that they are subsumed by laws such that from the presence of state S in circumstances C, together with the laws operating on the system of which it is a state in the circumstances, it follows that some other proposition is true. In this sense, the number of rings in the trunk of a tree carries information about its age, but it does not represent its age, and it does not have any representational or intentional states. This contrasts with the sense of ‘information’ on which

⁹ Why again? Why should TAG not just say: who cares about *original* intentionality, we just care about derived intentionality or as-if intentionality? The answer is that that would *trivialize* the proposal, and make the argumentation they actually offer *unnecessary*. This is why it would not do for TAG to adopt Dennett’s intentional stance account of when some system engages in cognition (Dennett 1981). According to Dennett, a system is a cognitive system provided that its behavior is well predicted by adopting the intentional stance with respect to it. Putting aside the problem that this characterization of when a system is a cognitive system presupposes an independent understanding of it (in presupposing an independent understanding of what it is to take up the intentional stance), this would make as-if intentionality sufficient for intentionality, and the argument of the paper could be much shorter.

the state that carries the information represents something as being the case. In this sense, we can speak also about misinformation or misleading or incorrect information. That is the hallmark of the capacity for genuine representation, that what is represented can fail to be so (with the usual caveats about necessary truths).¹⁰

With these clarifications in mind, let's examine these case studies in "group cognition".

3.1 Stigmergic Path Formation

The case of stigmergic path formation is represented as a case in which a group of walkers solves a problem without any members of the group thinking of themselves as engaged in solving the problem or even thinking of there being a problem for the group to solve. This is not then, in the ordinary sense, a matter of the group solving a problem. If we say that an individual solves a problem, we imply that the individual undertook to solve the problem intentionally and figured out a solution to it. 'Solving a problem' in this sense expresses an *essentially intentional action type*. You have to be intending to do it in order to do it. To intend to do it you have to represent the problem as a problem and have solving it as a goal. You might in a sense solve a problem accidentally, but only if you were intentionally engaged in trying to solve it and accidentally hit upon a solution. The same goes for a group solving a problem. In the ordinary sense, then, the group is not solving a problem at all. They are not doing anything together intentionally, let alone solving a problem together.

Is there some other sense in which they might be said to be solving a problem? Perhaps we could think of it being someone else's goal to develop a system of paths that minimizes the total path length, someone who fortuitously hits upon the idea of using a stigmergic mechanism. For example, we might say that it is the goal of a group of university planners, and that they have figured out how to adjust the visibility and decay rates of paths of walkers so as to achieve minimization of total path length. Then we can specify whose problem it is: the university planners. But we can also say in this case *who solved it*. The university

10 It can often appear in these discussions as if authors are moving back and forth between one or another thin sense of 'information' and a sense that involves representations, without realizing the difference. To trade on the nomic dependence sense of 'information', or the sense of 'information' at work in formal information theory, in order to conclude that groups do information processing in a sense that would imply that they are cognizers, would have, as Bertrand Russell put it, "all the advantages of theft over honest toil."

planners! They solved it by setting up conditions in which walkers would make the appropriate paths by a stigmergic mechanism. Thus, this does not advance the thesis under discussion. For the problem, in this case, is not solved, or even represented, by the walkers themselves. If a system does not solve a problem for itself, then it is not engaged in a process of cognition directed toward solving the problem. There is no group cognition at all in the walkers, not in the ordinary sense of them setting themselves a problem and solving it, or in any other sense – at least not just in virtue of their realizing a stigmergic mechanism for minimizing total path length.

The principle at work here is this: a system can be said to solve a problem only if it represents the problem as a problem for it and it undertakes steps to solve the problem intentionally. This principle is not satisfied by the system of walkers in stigmergic path formation, and so the system of walkers cannot be said in any sense to solve a problem. Someone else's using the walkers' activity for the purpose (by controlling the parameters that lead to different sorts of path formation) does not entail that the walkers represent themselves as engaged in problem solving or having any purpose as a group in each individually walking about the university campus.

What if the system of paths arises without the intervention of any external agents, but just because of the facts about degree of path visibility and the decay rate with the other factors? But now it is neither the goal of the walkers nor of anyone else, and so no problem has been presented to or solved by anyone. To see that, in this case, it is nonsense to speak of a problem being solved, notice that *for any path system that emerges, one could characterize some goal that a group may have that would be satisfied by that system of paths*. It hardly follows that no matter what they do, they have solved a problem. That would empty the notion of any content whatsoever.

What if as a matter of fact there is something *beneficial*, for the group of walkers, about the path system that arises, though no member of the group or any other was thinking about it? This does not make any difference. It is an accident that the result arrived at is beneficial. No one was aiming at it. You might as well say that if I take a walk for exercise and stumble on a chest of treasure, I solved the problem of how to pay off the mortgage on my house.

We may say: it is as if the system has solved a problem for itself! Yes, but the difference between as-if and real problem solving is whether or not one is representing the goal and trying to achieve it.

What if the fact that the result is beneficial in some way helps to explain why people in the circumstances have the particular proclivities that they do, to follow the paths of others to the extent that they do? It is quite plausible that an explanation of this type is available for stigmergic mechanisms in ant and

termite colonies and in other similar patterns of behavior among social animals in the natural world that arise from stigmergic mechanisms. It is not plausible that this is going on in the case of the walkers, of course. However, *even if it were*, the attractive feature of these sorts of explanations is that they show how apparently intelligent behavior can arise in the absence of having to attribute to organisms any representation of the goals or any intentions to pursue them. That is, they explain away the appearance that intelligence is involved. In fact, stigmergic mechanisms which yield benefits for a group in its environment via actions that are not intended to be coordinated with others provide a classic example of a mechanism for explaining how a group can act *as if* it were intelligently pursuing a goal when it is not.

What of the social parity principle? “If, *in confronting some task*, a group collectively functions in a process which, were it done in the head, would be accepted as a cognitive process, then that group is performing that cognitive process.” We will come back to this when we discuss the other examples, but in the present case it is sufficient to note that in this case the group does not *confront a task*, at least in the sense that it does not take up a task that it intends to solve. The social parity principle therefore has no application in this case.

3.2 The Collective Coloring Problem

In contrast to the as-if problem solving of stigmergic mechanisms, the collective coloring problem is a genuine group task, conceived and represented as such by the members of the group who engage in the task, as are the other examples we will consider. Perhaps here then we will find genuine group level cognitive processes.

Problem solving is an intentional activity. When a group solves a problem it does something. As a preliminary to our discussion, it will be useful to ask what the logical form of group action sentences is. Let’s put aside for a moment problem solving and concentrate on some other type of action. Consider [1] and [2].

[1] I sang the national anthem.

[2] We sang the national anthem.

On the standard event analysis, the logical form of [1] can be rendered in [1L]. (I will elide some details for the sake of simplicity – they will not affect our discussion; see (Ludwig 2007, 2010) and note 11.)

[1L] $(\exists e)(\text{agent}(e, \mathbf{I}) \text{ and singing}(e, \text{the national anthem}))^{11}$

Here ‘singing(e, the national anthem)’ is to be interpreted as expressing the type of sonic event that is brought about by singing but stripped of the implication that it is brought about by an agent. The agency implied by the action verb is separated out in the first conjunct. By analogy, we might think that the proper analysis of [2] is [2L].

[2L] $(\exists e)(\text{agent}(e, \mathbf{we}) \text{ and singing}(e, \text{the national anthem}))$

If this is correct, then perhaps we have an even more straightforward argument of genuine group level cognitive processes, for [2L] says that the group as such stands in the agency relation to the singing the national anthem, and so has intentions, desires, beliefs and so on. Of course, this seems absurd, as if we could never sing a duet because there would always be in addition to each of us as individual agents the two of us together as a third agent with its own goals and intentions.

And of course that is the right reaction, and even proponents of genuine group level problem solving seem to recognize that the argument cannot be that straightforward since they deny that aggregative or disjunctive problem solving amounts to genuine group cognition, though we can say in these cases correctly that the group solved the problem.

Fortunately, [2L] does not get the logical form right. [2] is ambiguous between a distributive and collective reading. On the distributive reading it means each of us sang the national anthem (e.g., as our audition piece), and we can represent this as in [2d]. Once we see that the noun phrase in this case is a distributive quantifier over members of the group, we can see that simply switching the scope of the quantifiers yields a representation whose truth conditions coincide with those we intuitively recognize as necessary and sufficient for the truth of [2] on the collective reading.

11 The full version includes a quantifier over a primitive action, a parameter on the agency relation set by the action verb, and the requirement that the subject be the only agent in the way required by the verb of the consequent event, and a quantifier over times. Omitting the quantifier over time: $(\exists e)((\exists f)[\text{agent}(f, \mathbf{I}) \ \& \ \text{by}_v(f, e) \ \text{and} \ (\text{only } y = \mathbf{I})(\exists f')(\text{agent}(f', y) \ \& \ \text{by}_v(f', e))] \ \& \ \text{singing}(e, \text{the national anthem})$. Here ‘agent(f, I)’ is interpreted as ‘I am a primitive agent of f’, and ‘by_v(f, e)’ is interpreted as ‘e is brought about by f v-ly’ where ‘v’ is a place holder for the form of agency required by the verb, in the case of ‘sing’, direct causation without the intervention of another agent, (I cannot sing the national anthem by paying you to do it). In the representation in the text, we can think of ‘agent(e, I)’ as abbreviating ‘ $(\exists f)[\text{agent}(f, \mathbf{I}) \ \& \ \text{by}_v(f, e) \ \text{and} \ (\text{only } y = \mathbf{I})(\exists f')(\text{agent}(f', y) \ \& \ \text{by}_v(f', e))]$ ’.

[2d] (Each x of us)($\exists e$)(agent(e, x) and singing(e , the national anthem))

[2c] ($\exists e$)(Each x of us)(agent(e, x) and singing(e , the national anthem))¹²

Joint action is simply a matter of there being multiple agents of an event (and so for joint intentional action is as well, though, in this case, the agents must intend that they do the action together in accordance with a shared plan at the time of acting (Ludwig 2007)).

Now, one might say: yes, that is so *in the case* in which the joint action is a matter of bringing about an event that is an aggregation of events of the same type that the group brings about, as in the case of singing the national anthem. But what about a case in which what the group brings about is not anything that any of the individuals do themselves? The answer is that this is not different at all. Thus, the distributive and collective readings of [3] are [3d] and [3c].

[3] We lifted a piano

[3d] (Each x of us)($\exists e$)(agent(e, x) and lifting(e , a piano))

[3c] ($\exists e$)(Each x of us)(agent(e, x) and lifting(e , a piano))

The consequent event here is the piano's going up ('lifting(e, x)' here, like 'singing(e, x)' above, is to be interpreted as expressing the type of event brought about by lifting (interpreted as a verb of agency) but to *strip off* the implication that it was brought about by an agent). So here lifting a piano is not something that each agent in the group does him or herself, but it is not something that requires a super-agent either. It is just a matter of each of them being an agent of the piano going up. The same point applies to a task in which there is a differentiation of roles, as when an orchestra performs a symphony.

¹² As represented here, it looks as if [2c] entails [2d], on the assumption that 'agent(e, x)' may relate more than one agent to the same event. And while this will not cause problems in this case, it would in cases in which the group does something that none of its members themselves do, as in the next example of lifting a piano together. However, as I pointed out in note 11, 'agent(e, x)' abbreviates a relation that requires there be one but at most one agent of e in the way required by the verb 'sing'. From each of us being an agent of one and the same event in a certain way it does not follow that each of us is the unique agent of an event in that way. But this does raise the question of how the uniqueness condition is understood on the collective reading in [2c], for obviously it cannot be that we require there is an event of which each of us is the unique agent. It is rather that we require that only members of the group (us) are agents of it in the right way. The full dress version of [2c] (omitting only a quantifier over time to handle tense) follows: ($\exists e$)(Each x of us)($\exists f$)[agent(f, x) & by_v(f, e) and (only y among us)($\exists f'$)(agent(f', y) & by_v(f', e))] & singing(e , the national anthem). See (Ludwig 2007) for a derivation of this.

The point of this is that the same thing goes for problem solving. The proper analysis of [4] is [4d] on the distributive and [4c] on the collective reading.

[4] We solved the crossword puzzle.

[4d] (Each x of us)($\exists e$)(agent(e, x) and solving($e, \text{the crossword puzzle}$))

[4c] ($\exists e$)(Each x of us)(agent(e, x) and solving($e, \text{the crossword puzzle}$))

Clearly, the truth of the claim that a group has solved a problem of some sort, or carried out a task that required some cognitive effort, implies only that its members contributed in some way or other to the result. It implies nothing about there being a group over and above the members who contributed something. We should not be fooled, then, into thinking that the mere fact that it is true to say that a group has solved a problem shows that there is a *group agent* that has solved the problem. For a group to solve a problem in the ordinary sense is for its members all intentionally to make a contribution to its solution.

In the collective coloring problem, the first thing to note is that in saying that the group has solved the problem, we mean no more than that each member of the group has made a contribution to solving it. One has to say more to show that there is a system other than the members of the group that can be said in its own right to be solving a problem. Thus, the observation that it is obvious that groups solve problems, even problems that their members cannot solve, and that they do so by a division of cognitive labor, does not show that the group is a problem solver in the sense in which individuals are. It does not show that the group has cognitive problem solving capacities in the sense in which individuals do. It does not show that the group has intentional states, other than those of its members. It does not show that the group per se cognizes, or reasons, or thinks, or has goals, and so on.

The main case for counting a group's solving the collective color problem as a group level cognitive process seems to be that *if the problems were solved by individual humans*, it would be interpreted as requiring relatively sophisticated problem solving strategies (see the quotation above from p. 386 in Section 2.3.2). The argument form seems to be the following:

1. A group working together did x .
2. If an individual had done x , then the individual would have had to have property F .
3. Therefore, a group as such has property F .

But now we can easily see that this argument is invalid. Consider:

1. A group working together lifted the piano.
2. If an individual had lifted the piano, then the individual would have been the agent of the lifting.
3. Therefore, a group as such was the agent of the lifting.

The premises may be true though the conclusion is false.

The argument just considered does not in fact invoke the social parity principle. Perhaps, however, it was intended to do so, despite what seems to be the form of argument in the text. Let us therefore consider the social parity principle in this connection.

If, in confronting some task, a group collectively functions in a process which, were it done in the head, would be accepted as a cognitive process, then that group is performing that cognitive process.

In the collective coloring problem, a group is genuinely confronting a task, unlike the case of stigmergic path formation. But what is the process in this case?

First, just to clear the ground, let's ask whether what the group does is like what an individual would do in solving the problem. If it were, then it would be straightforward to describe the group as such as engaged in a cognitive process. Suppose that an individual does solve the coloring problem in the cycle topology. This actually is fairly simple. We can see pretty much immediately any number of solutions, and among the simplest is simply alternating between two colors around the circle. The process, say in a particular case, involves framing the problem, representing the topology, recognizing that one color flanked by another color on either side meets the condition for a single node to have differently colored neighbors, and then that either of them would satisfy it as well if they were flanked by nodes of the same color as the one between them (assuming an even number of nodes), and so on. However, when the group solves the problem, is that what the group does? That is, does the group itself represent the topology, recognize that one color flanked by the same color on either side meets the condition on a single node to have differently colored neighbors, and so on? The basic description of what happens does not include this, and it is not necessary in order for us to understand how the group solves the problem. To assume this begs the question.

Next, let's consider the application of the social parity principle, which asks us to consider how we regard the process if it were to go on in the head. What is the process? The process the group goes through is a matter of different agents making decisions in the light, say, of what they can tell of their immediate

neighbors, until they all settle on a stable pattern. But what would it be for that process to occur in an individual? Are we to imagine individuals who are actually collections of agents working independently? That is not what actually happens in us when we solve such a problem. But *if that did happen “in the head”*, why would that show that the *agent* in whose head this bizarre process was occurring had solved a problem?¹³ He might not be thinking about the problem at all. However, let us say we abstract from the fact the process is implemented by agents. It is still unclear how we are to think of it as implemented in the head, and how it connects with the agent’s attitudes. Suppose, however, that we somehow clear this up. Still, what reason would we have for thinking such a process occurring in the head would be a cognitive process? Apart from the obscurity of the idea that the same process is going on in the head, the basic problem is that *unless there is a reason already to think it is in the relevant sense a cognitive process, thinking of it as occurring in the head gives us no additional purchase on whether it is*.

One might respond that we should think of the process not in terms of the details of how it is carried out, but simply in terms of what its input and output is. Then the principle is clarified as follows:

If, in confronting some task, a group collectively functions in a process that has as input X and as output Y, and if a process done in the head which has input X and output Y is a cognitive process, then that group is performing that cognitive process.

However, this principle accepts the crudest form of behaviorism about cognition, namely, that all that matters is input and output, not the processing in between. The social parity principle, so interpreted, would yield the desired conclusion in application to these cases, but the principle so interpreted is false (Block 1981).¹⁴

13 A referee suggested that to implement the process in an individual, we should imagine that the individual sequentially occupies the position of each member of the group. Would this be the same process? The individual of course represents himself as solving the problem by sequentially occupying each of the nodes. Is that the process that the group was engaged in? That is, was the group per se (as opposed to its members) representing itself as solving the problem by sequentially occupying each of the nodes? No. Well, but is that not an *uncharitable* way of taking the suggestion? Not if one of the reasons we take the individual to be engaged in a cognitive activity is that he is representing himself as solving the problem by adopting a certain strategy. Apart from this, it is not clear that one could replicate the process in this case by sequentially occupying the nodes. That would mean that each node occupier could only get information about changes in at most one node at a time. But there is no such restriction on information in the collective coloring problem.

14 Thinking about Block’s example of the list look-up program strategy for passing the Turing Test (1981) helps to see the problem with invoking the social parity principle. Block’s point is that the list look-up program is clearly not an instance of intelligence in the system that instanti-

The difficulty we have run into is this. We can interpret the social parity principle either in a thick sense or a thin sense. On the thick interpretation, for the same process to occur in a group and an individual, the details of the process matter. While it is unclear in these applications what details are supposed to matter, no matter how we adjudicate this, the basic problem is that unless we already have a reason to think that the process described at the relevant level of detail is a cognitive process, thinking of it as going on in the head is not going to help. If we knew the process was a sequence of intentional states obeying norms of reasoning, then if it were going on in the head, we would count it as a cognitive process. But then we do not need to think of it as in the head to recognize it as a cognitive process. But if it is not obvious already, how does thinking of it as going on in the head help? On the thick interpretation, the principle is useless. On the thin interpretation, we move from the claim that some cognitive process in an individual has a certain input and a certain output and the claim that a process in a group solving a problem has the same input and output, to the claim the group is engaged in a cognitive process. This would get the right result on the basis of claims about the group that no one would dispute. In application to the case of the collective coloring problem it would entail that the group itself was a subject of cognition. However, the principle is false for the same reason that behaviorism is false: inner processing matters, in addition to the input-output function, to whether a system is displaying intelligence. In sum, if we interpret the principle so that it is guaranteed to get the right result (the thin interpretation), it is false, while if we interpret it in a way that will secure its truth (the thick interpretation), it gives us no additional purchase on whether the process going on in the group is a cognitive process.¹⁵

ates it, though from the outside it looks as if it is. We can imagine that the program is actually implemented by a very large team of individual agents through a cognitive division of labor. Then we have a system that is implemented by a group of agents which acts as if it is intelligent. Now we can try to apply the social parity principle: If, in confronting some task, a group collectively functions in a process which, were it done in the head, would be accepted as a cognitive process, then that group is performing that cognitive process. Ask yourself how this helps in the present application. If we can see that it is not sufficient for the system (as opposed to its parts) to be intelligent, then obviously if it were done in the head, it would not be accepted as a cognitive process. That is what we could say about the list look-up program instantiated by a group of agents. If we can see that it is sufficient, then if it were done in the head, it would be sufficient as well, but then the principle does not help to establish the relevant claim. But suppose there is a question about it. Then the question arises as well if it is done inside the head of an individual – unless we bite the bullet and accept a thorough going behaviorism. But that is simply untenable.

15 The attraction of the social parity principle is that it seems to provide us with a means of settling a question without actually doing the hard work of saying what is required for cognition. But there is not really any way of avoiding that. The only sort of process in a group which, if done

What about the markers of group cognition R1–R3? These are supposed to be marks of genuine group cognition, but they identify necessary, not sufficient conditions for group cognition. The fact that a group’s solving a problem is not merely a matter of aggregating their contributions (e.g., because they interact with each other) does not show that there is a group level cognizer (think of the symphony orchestra again or a team wielding a two man saw – non-aggregative agency is not ipso facto group level agency). The fact that solving a problem involves a differentiation of roles does not entail that there is a group level cognizer. Even the fact that there is some organizational level fact that helps explain how the group does what it does is neutral with respect to whether there is a group cognizer on the scene, because the question is whether the group level organization feature is properly cognitive.

The only principles that pertain to sufficient conditions are M1 and M2. But M1, the social parity principle, is no help on the thick interpretation and is false on the thin interpretation. M2 does not apply in the present case because it seems clear that the individuals in the collective coloring problem are able to solve the problem themselves.

3.3 Acquired Division of Labor

With this discussion in hand, we can dispatch the other cases fairly quickly. In the case of the acquired division of labor, involving the “Group Binary Search” game, we have individuals who sort themselves into different roles as reactive or nonreactive, through multiple rounds of the game, which improves their performance. It cannot be said here that they do better than individuals, but they perform the same task as individuals and they increase efficiency by differentiating their roles. We might also say that the group learns better how to do the task set them, perhaps even without setting themselves to do so.

Again: the fact that they solve the problem, or carry out the task, does not show that the group per se is an agent of the solution. But there are group level cognitive processes in the sense we are interested in only if the group per se is an agent of the solution. For what a group cognitive level process of the sort that we

in the head, would count as a cognitive process, is a cognitive process. So to know whether a process in a group, if done in the head, would count as a cognitive process, we need to know whether the process in the group is a cognitive process. Being asked to consider it going on in the head is not going to help with this, except insofar as it jars us into recognizing properties of it as cognitive which we could have recognized independently, dislodging a kind of prejudice. The only way around this is to adopt the input-output function reading of the principle on which it is false.

are interested in requires is that there be group level intentional states, a group level thinker, or cognizer, a group level possessor of representations of the task, a group level desire and intention to pursue it, and group level beliefs about how to do it, and, hence, a group level agent. Nothing follows, as we have seen, about there being a group level agent from the fact that the group solves the problem, because all this comes to is saying that each member of the group intentionally contributes to its solution.

What about the fact that the group *learns* how to do the job better? This does not require anything other than a distributive reading: each member of the group learns how he can contribute better to what they do in light of how others modify their contributions. What about the suggestion that they get better at it, but by way of a mechanism that does not involve their individually trying to get better at it? If we take that seriously, then we should also say that though they got better at it, they did not *learn* how to do it better.

3.4 Group Remembering and Transactive Memory Systems

The final case study is group remembering and transactive memory systems. A transactive memory system is a group of individuals who perform a task drawing on diverse skills among its members, where the knowledge and skills required for the task are distributed across the members of the group, and where those skills also include how and when to prompt other members in order to elicit an appropriate contribution. Suppose the task that the group performs is the assembly of a radio. Then we can say that though no member of the group knows how to assemble a radio, the group does, and we can say that though no member of the group remembers all the steps, the group as a whole does. Here we have not merely a division of labor or even a division of cognitive labor, but also a division of skills and knowledge relevant to the task. And here, as well, the task, if done by an individual, would be counted as the expression of various cognitive capacities and skills that would be taken to be the expression of intelligence, the possession of knowledge, and the exercise of memory. Surely in this case we must say that the group itself possesses cognitive capacities, remembers, has skills, knows things, and so on.

Not at all! First, what the group does is to assemble a radio. This is a matter of each of them making a contribution to bringing it about that a radio is assembled. This does not require, and our ordinary ways of speaking do not entail, that there is a super-agent distinct from any of the agents in the group who do the assembling. Second, we have a complete explanation of how the contributions each makes to the outcome result in it by appeal to their differing skills, knowledge, and so on, and the interactive effects between them (prompting, eliciting and

so on). No group level *cognitive* processes are required. They are explanatorily redundant. Third, though the group is able to do something that the individuals who are members of it cannot do alone, what this comes to is that it is possible for the group to do something by their working together, exercising in coordination their individual capacities in making contributions to the outcome. It does not follow from this that the group has a skill in the sense that an individual would who was able to perform the same task. In particular, it does not entail that the group as such represents all the components of the task, or as such has the manual skills required for assembly tasks at various points. These are all expressed through individuals.

Consider M2 in connection with this. M2 states that: We have evidence for group level cognition when the capacities for problem solving that can be ascribed to the group are distinct from those possessed by any individual members of the group. If we interpret “capacities for problem solving” to mean just that the group can solve a problem that its individual members cannot, then the principle is false. It is false because all this comes to is that an outcome can be achieved by a group that its individual members cannot bring about alone. Consider a complex problem that is solved in the following way. One person works on it and gets the problem to a certain point. Someone else then takes up the problem from that point and works on it some more. Then someone else takes it up and finishes. Suppose that none of them could have done it alone, and in fact suppose each of them devoted a lifetime to it and the next person built on the previous person’s work. They solved the problem. But it could hardly follow that there was a group level cognitive process because the members of the group that solved the problem did not exist at the same time.¹⁶ The only cognitive processes going on at any time were in the heads of individuals. On the other hand, if “capacities for problem solving” is interpreted to mean the capacities to undergo *group level cognitive processes*, then we can grant the principle is true, but that the capacities are different from those of the individuals is neither here nor there, and the principle cannot be applied in the present case without begging the question.

Fourth, the application of the social parity principle is as flawed in this case as it is in the case of its application to the collective coloring problem. M1 says:

¹⁶ The objection of course is not that cognitive processes cannot have any duration, but that if the group as such is supposed to be the subject of an intentional state at a time t , the group must exist at t . From the example, we can see that different agents making complementary contributions to solving a problem that none of them alone can solve cannot entail that there is a group that has intentional states per se because this requires that the group exist at the time it has the intentional state and that requires that its members all coexist at some time at least – but in the example, while they all contribute to solving the problem, they never coexist together.

“If, in confronting some task, a group collectively functions in a process which, were it done in the head, would be accepted as a cognitive process, then that group is performing that cognitive process.” The trouble is again that if, on the one hand, we interpret “process” simply in terms of the output given the input (the thin interpretation), we can grant that what an individual does who performs the cognitive components of assembling a radio is the same process as the group performs. But if this is how we interpret the principle, then it assumes a form of behaviorism about cognition, namely, that cognition is simply an input-output function and the internal processing does not matter. But that has long been recognized to be untenable. But if, on the other hand, we interpret “process” in a way that requires us to imagine that what the group does itself be done in the head (the thick interpretation), then unless we already have reason to think the process is cognitive process, this is not going to help. The group does it by dividing up the cognitive and other labor between individual agents. That is not in fact the way an individual would do it “in the head”. And if what the group does occurred “in the head” (however we are supposed to imagine that), it is no more obvious that the individual agent would be solving the problem because that process occurred “in the head” than it is when it occurs “outside the head.”

Yet, surely the TMS is a *memory system*. As TAG say, in a TMS, “(A) information is stored from experiences, (B) [there is] subsequent access to this stored information, and (C) subsequent use of the access information to deal with related situation more effectively” (p. 389). We may so-define a memory system, but the information is stored in individual memories, it is accessed by individuals and shared through individual actions, and the use of it made is by individuals, each contributing to something none of them bring about by their efforts alone. There is no group cognition here in any sense other than that the group carries out a task together that draws on complementary skills.

4 Objections and Replies

Objection 1. You are denying that groups can genuinely solve problems because they have agents as their constituents. But suppose it turned out that we had agents as constituents. We would not deny that we were able to solve problems or had thoughts.

Reply: I am not saying that groups cannot genuinely solve problems because they have agents as their constituents. I am denying that any of the examples of groups solving problems have been shown to be examples of group level cognitive processes.

Objection 2. But you must be denying functionalism, because all that is needed for the in principle possibility of group level cognitive processes is that functionalism is true.

Reply: It is true that functionalism is all that is needed for the in principle possibility (emergentism too allows this). But I am not denying functionalism. The case was not being made on the basis of a functional analysis of cognition and intentionality, but on other grounds, and those grounds were found wanting. It is, furthermore, completely implausible to suggest that in the cases we have considered the groups having anything like the functional organization of individuals who could solve the same problems.

Objection 3. You have to admit that groups have different capacities for solving problems than individuals. Groups can sometimes solve problems that none of their individual members can and they can sometimes solve problems better, more quickly, or more reliably than individuals. That shows by itself that they have cognitive capacities that are distinct from the capacities of the individuals who make them up.

Reply: We have been over this ground already. If we mean by ‘group G has the capacity to solve problem S’ that G can solve the problem, then this just means that their multiple contributions can determine a solution. But this does not entail that the capacity is a capacity of the group per se to represent the problem, or its solution, or that the group per se has intentional states and so on. But if we mean by ‘group G has the capacity to solve problem S’ that G has cognitive states, represents the problems, and the solution, then I do not have to admit that groups have capacities for solving problems that their members do not. That is just begging the question. What was wanted was an argument for that conclusion.

Objection 4. Well, all that is required for the thesis is that groups have cognitive capacities, and undergo group level cognitive processes, in the sense that groups do solve problems by a division of labor where the individual members of the group do not have the skills, knowledge, or ability individually to solve the problem. Surely the proponents of the view never intended to be claiming that there were group level cognitive processes in your sense!

Reply: Well, if that is what is meant, then it can be granted that groups have cognitive capacities, and undergo cognitive processes, and so on. But now the thesis is utterly trivial, and it is uninteresting, and it is only in a Pickwickian sense that we can say that the examples show that group level cognition takes place.

Objection 5. Why think there are only the two options: full-blooded intentional states and reasoning, on the one hand, and behaviorism on the other? Surely, there is some intermediate position that TAG could occupy? Surely that is what they intended to occupy? As they say (see also note 7):

... we confine our attention to specific cognitive capacities that have traditionally been studied by psychologists at the level of individuals, but have more recently been used to analyze the behavior of groups as a whole. ... [we] seek ... successful characterizations of group behavior that invoke psychological constructs at the level of groups. (p. 385).

Isn't the idea to invoke "psychological constructs", that is, causal-structural models of cognition that are actually invoked in cognitive psychology at the level of individuals at the level of groups, but not necessarily ordinary notions? Isn't that what is really going on? And, if so, does not that show that these criticisms miss the mark, and that the project manages to steer between embracing behaviorism, on the one hand, and wildly implausible argumentative leaps from rather mundane observations about groups solving problems to there being group level intentional states, on the other?

Reply: First (see also note 7) TAG say, by way of follow up: "In this paper, we focus on group problem solving ... and on group memory" (p. 385). Both of these involve intentional, representational states. If these are the examples of "psychological constructs", they are just garden-variety cognition. Furthermore, in their concluding section, they repeat their general orientation:

... we suggest to break down the indistinct "ability to cognize" into a number of fairly well-understood capacities such as memory, attention, learning and problem solving. It makes good sense to call these capacities cognitive, based on a firm but defeasible *pre-theoretic understanding* of the distinctive roles which they play in the production of intelligent forms of behavior. (p. 393; emphasis added)

They have in mind memory, attention, learning, and problem solving. These are everyday mental capacities and processes. That's why there is not supposed to be any controversy about whether if groups have memory, pay attention, learn, and solve problems, there is group level cognition. It seems redundant to add that they should be considered cognitive capacities based on our pretheoretic understanding of the roles they play in the production of intelligent behavior, since it is not controversial, but this does show it is the pre-theoretic understanding of them that is at issue as well. On the face it – and it is in the light of this that the thesis appears both startling and significant – they aim to show that groups *themselves* are problem solvers, have memory, learn, pay attention, and so on. But they have

not shown this. (What significance should we assign to the modifier ‘defeasible’? This attaches to the roles memory, attention, etc., play in production of intelligent forms of behavior, so the idea is presumably that we may find more about their roles through empirical investigation than present in our pre-theoretic understanding of them.) Second, suppose we ask what cognitive models are supposed to be shown to be present in each of the examples? What about the case of path formation by stigmergic mechanisms? The only appeal in this case is to the idea that the group solves a problem. The same goes for the collective coloring problem and the example of acquired division of labor. No independently articulated psychological construct is introduced, shown to apply at the level of individuals, and then shown to apply at the level of groups. The thought that the TMS is a *memory system* we have dealt with already. All the conditions specified come to in the case of the group is that they did something together drawing on their individual memories and skills in a way that is complementary. To *relabel* this as group level cognition (effectively stipulating one will use ‘group level cognition’ to mean *distributed cognition which involves individuals contributing in complementary ways*) is not to advance a thesis of any substance.

5 Conclusion

Here are the morals of the discussion:

1. Stigmergic mechanisms that give rise to beneficial results for a population, even when their presence can be explained by the benefit, are not instances of group level problem solving, because they are not instances of problem solving at all. Genuine problem solving requires representing the problem and the goal of solving it and efforts directed at a solution. None of this is present in the operation of stigmergic mechanisms, and, in fact, the beauty of explanation by stigmergic mechanisms is that none of that is required.
2. The social parity principle has a thin behaviorist reading on which sameness of process come to preserving the input-output function, and a thick process reading on which we must preserve significant features of the process connecting input and output. The behaviorist reading will secure in the cases examined that there is a process in the examples discussed that goes on in a group that is the same as in an individual, but on this reading, the principle is false precisely because it presupposes a false theory of mental concepts, namely, behaviorism. On the thick reading, the principle cannot give us any additional purchase on whether the process in the group is a cognitive process.

3. The principle that group level capacities for problem solving differing from individual capacities provides evidence for emergent group level cognition has a thin and thick reading as well. On the thin reading, the notion of a problem solving capacity is cashed out in terms of it being true that the group solves the problem. On the thick reading, we understand the capacities to be the same as the capacities would be in an individual who solves the problem intentionally. On the thin reading, the principle is false. On the thick reading it cannot be applied without begging the question.
4. The various necessary conditions on emergent cognition are neither here nor there for those conditions may all be met and yet a system not have any group level cognitive states, processes or capacities.
5. Group problem solving in ordinary cases is simply a matter of a group setting itself the goal of solving a problem and undertaking intentionally to pursue it, where all members of the group make contributions to solving the problem they set themselves. Whether this involves a division of cognitive labor and cognitive skills and knowledge or not, it is nothing more than there being multiple agents of the production of a solution to the problem. No super agent is needed. None is present.
6. The only sense in which we may allow that there is group level cognitive activity is given simply by the description of what occurs, and that gives us a Pickwickian sense of group level cognitive activity, both trivial and uninteresting.

If these arguments so obviously fail to hold up to scrutiny, why has it seemed to anyone that we can get genuine group level cognition so cheaply? Part of it may be what I called terminological drift or sliding (see note 3), using the terms ‘cognitive system’, ‘cognitive process’, ‘cognitive property’, and the like, in a rather anodyne sense, and sliding to a more robust interpretation, without realizing it. But I think a larger part of it comes from confusion about the import of the surface form of our attributions of problem solving to groups. We say that groups have solved complex problems, when we also know that none of the individuals would have been able to do it on their own. Since solving a problem is an intentional activity, and it requires representing a goal and working out how to arrive at it, it seems that we are forced to think of there being an intelligent agent who does the solving who is not identical with any members of the group (a locus for that intelligence distinct from any member), that is, a group agent who solves the problem. The rest of the apparatus is introduced because we can see that there are cases in which we would say a group solved a problem and speak truly but in which it just seems completely implausible that there is a group agent that solved a problem, as when all the members of a team work independently and adopt the best solution to present as theirs, or in which they aggregate judgments and rely on the

Condorcet Jury theorem. But what drives the view is the idea that if some result is achieved that cannot be achieved by individual agents alone and is intelligent goal directed behavior, it calls for a group cognizer to be the locus of the problem solving. Once this is brought out into the open, the fallacy becomes obvious. Just as four people lifting a piano does not require a super-agent with huge biceps and thighs, but only multiple agents of a piano going up, so solving a problem in a group by a division of labor does not require a super-agent with an amazing intellect but only multiple agents contributing in multiple ways to bringing about a result. The only representations of the goals are in the heads of individual agents in the group, and the only representations of the solution are in the heads of individual agents in the group. And if none of them in fact represent the whole solution, the representation is distributed across members of the group, as a house design may be distributed across blueprints for the different floors, but it is still not in any super-mind.

Bibliography

- Adams, F. R. and K. Aizawa (2008): *The Bounds of Cognition*. Oxford: Blackwell.
- Block, N. (1981): "Psychologism and Behaviorism". In: *Philosophical Review* 90, p. 5–43.
- Button, G. (2008): "Against 'Distributed Cognition'". In: *Theory, Culture and Society* 25. No. 2, p. 87–104.
- Clark, A. (2009): "Reply to Fodor". In: *London Review of Books* 31. No. 6.
- Dennett, D. C. (1981): "True believers: The intentional strategy and why it works". In: A. F. Heath (Ed.): *Scientific Explanation: Papers Based on Herbert Spencer Lectures Given in the University of Oxford*. Clarendon, p. 150–167.
- Fodor, J. (2009): "Where is My Mind?" In: *London Review of Books* 31. No. 3, p. 13–15.
- French, P. (1979): "The Corporation as a Moral Person". In: *American Philosophical Quarterly* 16. No. 3, p. 207–215.
- Gallagher, S. (2013): "The Socially Extended Mind". In: *Cognitive Systems Research* 25–26, 4–12. doi: <http://dx.doi.org/10.1016/j.cogsys.2013.03.008>.
- Gilbert, M. (1989): *On Social Facts*. London: Routledge.
- Goldstone, R. L. and M. E. Roberts (2006): "Self-organized Trail Systems in Groups of Humans". In: *Complexity* 11, p. 43–50.
- Goldstone, R. L., M. E. Roberts, and T. M. Gureckis (2008): "Emergent Processes in Group Behavior". In: *Current Directions in Psychological Science* 17. No. 1, p. 10–15.
- Grassé, P.-P. (1959): "La Reconstruction du Nid e les Coordinations inter-individuelles chez *Bellicositermes natalensis* et *Cubitermes* sp. La Theorie de la stigmergie: Essai d'interpretation des termites constructeures". In: *Insectes Sociaux* 6, p. 41–83.
- Grofman, B., G. Owen, and S. L. Feld (1983): "Thirteen Theorems in Search of the Truth". In: *Theory and Decision* 15, p. 261–278.
- Gureckis, T. M. and R. L. Goldstone (2006): "Thinking in Groups". In: *Pragmatics & Cognition* 14, p. 293–311.

- Helbing, D., J. Keltsch, and P. Molinár (1997a): “Modelling the Evolution of Human Trail Systems”. *Nature and System* 388, p. 47–50.
- Helbing, D., F. Schweitzer, J. Keltsch, and P. Molnár (1997b): “Active Walker Model for the Formation of Human and Animal Trail Systems”. In: *Physical Review E* 56, p. 2527–2539.
- Heylighen, F., M. Heath, and F. Van Oerwalle (2004): *The Emergence of Distributed Cognition: A Conceptual Framework*. Paper presented at the Collective Intentionality IV, Siena, Italy.
- Hutchins, E. (1995): *Cognition in the Wild*. Cambridge, Mass.: MIT Press.
- Laughlin, P. R. (2011): *Group Problem Solving*. Princeton: Princeton University Press.
- List, C. (2003): “Distributed Cognition: A Perspective from Social Choice Theory”. In: M. Albert, D. Schmidtchen, and S. Voigt (Eds.): *Scientific Competition: Theory and Policy, Conferences on New Political Economy*. Mohr Siebeck.
- List, C. and P. Pettit (2011): *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- Ludwig, K. (2007): “Collective Intentional Behavior from the Standpoint of Semantics”. In: *Noûs* 41. No. 3, p. 355–393.
- Ludwig, K. (2010): “Adverbs of Action”. In: T. O’Connor and C. Sandis (Eds.): *Blackwell Companion to the Philosophy of Action*. Oxford: Wiley-Blackwell.
- Michaelian, K. and J. Sutton (2013): “Distributed Cognition and Memory Research: History and Current Directions”. *Review of Philosophy and Psychology* 4. No. 1, p. 1–24.
- Moussaïd, M., S. Garnier, G. Theraulaz, and D. Helbing (2009): “Collective Information Processing and Pattern Formation in Swarms, Flocks, and Crowds”. In: *Topics in Cognitive Science* 1, p. 469–497.
- Page, S. E. (2007): *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton: Princeton University Press.
- Pettit, P. (2003): “Groups with Minds of their Own”. In: F. Schmitt (Ed.): *Socializing Metaphysics*. Rowman and Littlefield, p. 167–193.
- Pfeifer, J. and S. Sarkar (2006): *The Philosophy of Science: An Encyclopedia*. Psychology Press.
- Steiner, I. D. (1966): “Models for Inferring Relationships between Group Size and Potential Group Productivity”. In: *Behavior Science* 11, p. 273–283.
- Sutton, J., C. B. Harris, P. G. Keil, and A. J. Barnier (2010): “The Psychology of Memory, Extended Cognition, and Socially Distributed Remembering”. In: *Phenomenology and the Cognitive Sciences* 9. No. 4, p. 521–560.
- Szanto, T. (2014): “How to Share a Mind: Reconsidering the Group Mind Thesis”. In: *Phenomenology and the Cognitive Sciences* 13. No. 1, p. 99–120.
- Theiner, G. (2008): *From Extended Minds to Group Minds: Rethinking the Boundaries of the Mental*. (Ph.D.), Indiana University, Bloomington, Indiana.
- Theiner, G. (2009): “Making Sense of Group Cognition: The Curious Case of Transactive Memory Systems”. In: W. Christensen, E. Schier, and J. Sutton (Eds.): *ASCS09: Proceedings of the 9th Conference of the Australasian Society for Cognitive Science*. Macquarie Center for Cognitive Science, p. 334–342.
- Theiner, G. (2013a): “Onwards and Upwards with the Extended Mind: From Individual to Collective Epistemic Action”. In: L. Caporael, J. Griesemer, and W. Wimsatt (Eds.): *Developing Scaffolds*. MIT Press, p. 191–208.
- Theiner, G. (2013b): “Transactive Memory Systems: A Mechanistic Analysis of Emergent Group Memory”. In: *Review of Philosophy and Psychology* 4. No. 1, p. 65–89.

- Theiner, G. (2014): "A Beginner's Guide to Group Minds". In: J. Kallestrup and M. Sprevak (Eds.): *New Waves in Philosophy of Mind*. Palgrave Macmillan.
- Theiner, G. and T. O'Connor (2010): "The Emergence of Group Cognition". In: A. Corradini and T. O'Connor (Eds.): *Emergence in Science and Philosophy*. New York: Routledge.
- Theiner, G., C. Allen, and R. L. Goldstone (2010): "Recognizing Group Cognition". In: *Cognitive Systems Research* 11. No. 4, p. 378–395.
- Tollefsen, D. (2004): "Collective Epistemic Agency". In: *Southwest Philosophy Review* 20. No. 1, p. 55–66.
- Tollefsen, D. (2006): "From Extended Mind to Collective Mind". In: *Cognitive Systems Research* 7, p. 140–150.
- Wegner, D. M. (1986): "Transactive Memory: A Contemporary Analysis of the Group Mind". In: B. Mullen and G. R. Georthals (Eds.): *Theories of Group Behavior*. New York: Springer, p. 185–208.