


Provided for non-commercial research and education use.  
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



ELSEVIER

Available online at [www.sciencedirect.com](http://www.sciencedirect.com)

Cognitive Systems Research 11 (2010) 378–395

[www.elsevier.com/locate/cogsys](http://www.elsevier.com/locate/cogsys)

# Recognizing group cognition

Action editor: Leslie Marsh

Georg Theiner<sup>a,\*</sup>, Colin Allen<sup>b</sup>, Robert L. Goldstone<sup>c</sup><sup>a</sup> *Department of Philosophy, University of Alberta, Edmonton, Canada*<sup>b</sup> *Department of History and Philosophy of Science & Cognitive Science Program, Indiana University, Bloomington, IN, USA*<sup>c</sup> *Psychological and Brain Sciences Department & Cognitive Science Program, Indiana University, Bloomington, IN, USA*

Received 1 April 2010; accepted 5 May 2010

Available online 29 July 2010

## Abstract

In this paper, we approach the idea of group cognition from the perspective of the “extended mind” thesis, as a special case of the more general claim that systems larger than the individual human, but containing that human, are capable of cognition (Clark, 2008; Clark & Chalmers, 1998). Instead of deliberating about “the mark of the cognitive” (Adams & Aizawa, 2008), our discussion of group cognition is tied to particular cognitive capacities. We review recent studies of group problem solving and group memory which reveal that specific cognitive capacities that are commonly ascribed to individuals are also aptly ascribed at the level of groups. These case studies show how dense interactions among people within a group lead to both similarity-inducing and differentiating dynamics that affect the group’s ability to solve problems. This supports our claim that groups have organization-dependent cognitive capacities that go beyond the simple aggregation of the cognitive capacities of individuals. Group cognition is thus an emergent phenomenon in the sense of Wimsatt (1986). We further argue that anybody who rejects our strategy for showing that cognitive properties can be instantiated at multiple levels in the organizational hierarchy on *a priori* grounds is a “demergentist,” and thus incurs the burden of proof for explaining why cognitive properties are “stuck” at a certain level of organizational structure. Finally, we show that our analysis of group cognition escapes the “coupling-constitution” charge that has been leveled against the extended mind thesis (Adams & Aizawa, 2008).

© 2010 Elsevier B.V. All rights reserved.

**Keywords:** Group cognition; Distributed cognition; Extended mind; Emergence; Collective memory; Collective problem solving

## 1. Introduction

As thinking entities all too aware of our own cognition, it is natural for us to ignore or dismiss the possibility that larger systems which contain us among their proper parts can also be cognitive systems. Nonetheless, we maintain that cases of group cognition are empirically observed. In particular, we shall argue that groups of people can manifest cognitive capacities that go beyond the simple aggregation of the cognitive capacities of their individual members.

In asserting that groups cognize, we attribute mental properties to groups that are normally attributed to individuals. What are those properties? From a folk-psychological

perspective, candidate properties for group cognition might include intent, thought, intelligence, and consciousness. Going along with these might be other properties more typically studied directly by cognitive scientists such as memory, concepts, attention and learning. From a more abstract perspective one might try to characterize cognition using notions such as computation, information processing, generation and use of internal and external representations, and problem solving flexibility.

Many philosophers have strong intuitions that cognitive or mental properties should be attributed at exactly one level of organization: the individual organism. For some properties, we share these intuitions. It is very hard, for example, to imagine phenomenological consciousness at a group level. What is it like to be a group? Such failure of imagination is not an argument, of course, but it does

\* Corresponding author.

E-mail address: [georg.theiner@ualberta.ca](mailto:georg.theiner@ualberta.ca) (G. Theiner).

suggest that one should be careful to examine the sources of intuitions about the possibility of group minds or group cognition. 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 & 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. Further below, we illustrate this approach by discussing studies of group problem solving, and research on group memory.

Underlying our methodology is the idea that 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. Although in this paper we focus on cognition in groups of humans, we think of this as a more widespread natural phenomenon, illustrated, for example, by the fact that better anti-predator vigilance in mixed species flocks of birds is not simply the aggregation of many individual acts of vigilance, but due to specialization within the flock to the point where certain species are recognized by biologists as “sentinel species.” These sentinel species in turn benefit from the insect-flushing activities of members of other bird species comprising the flock (Munn, 1986). The flock solves the problem of feeding while remaining vigilant by dividing the cognitive labor of finding insects from the cognitive labor of detecting predators.

We know that many readers will not be convinced at this point that such a flock of birds is a problem solving system, still less that the flock as whole has its own cognitive properties. Hence the goal of our paper is to say more clearly why the idea of group cognition is neither trivial nor shrouded in metaphysical mystery, but ought to be taken seriously. In the pursuit of this goal, we will not take on those who think it is somehow impossible for groups of organisms (or other extensions of individual organisms) to possess cognitive properties, or that it is a category error to think that they can. (It should be possible to extrapolate our response to these views from what we say below.) Rather, we focus on responding to those critics who allow that socially or environmentally extended entities containing one or more organisms among their proper parts *could perhaps* possess cognitive properties distinct from those of

the individual organismic part(s), but who go on to deny that any of the cases offered by proponents of extended cognition *actually do* possess cognitive properties at the higher level. Such critics have followed a strategy of attempting to show that all the putative examples of extended cognition given to date can be analyzed into organisms possessing cognitive properties interacting through non-cognitive media. Given such an analysis, these critics claim, to attribute cognitive properties to the larger entity is to commit a compositional fallacy.

Our discussion is organized into the following sections. In Section 2, which we have divided into three subsections, we state our basic account of group cognition. In Section 2.1, we introduce the “extended mind” thesis, and motivate our approach of treating group cognition as a special case of the more general claim that cognition is not confined to the boundaries of the individual. In Section 2.2, we distinguish three ways in which group cognition can be said to be “emergent” with respect to the cognitive abilities of individuals. In Section 2.3, we introduce and discuss the inverse notion of “demergence,” and use it to ascertain how the burden of proof is appropriately distributed among proponents and opponents of group cognition. In Section 3, we make a detailed empirical case for group cognition in the area of collective problem solving investigated by Goldstone and colleagues, and the theory of transactive memory systems. Taking a closer look at these ongoing research programs will help us to highlight the explanatory virtues of understanding groups as cognitive systems. In Section 4, we show that our account of group cognition escapes the “coupling-constitution” charge that has been leveled against the extended mind thesis at large (Adams & Aizawa, 2008). We conclude our discussion with a brief reflection on the relationship between our account of group cognition and the scope of the “extended mind” thesis.

## 2. Towards an account of group cognition

### 2.1. Extending the “extended mind” to groups

Our claim that groups cognize should be seen a special case of the more general claim that systems larger than the individual human, but containing that human, are capable of cognition (Clark, 2003; Clark & Chalmers, 1998; Wilson, 2004). For instance, human beings are uniquely talented in modifying their environments through the creation and manipulation of “cognitive artifacts” (Norman, 1991) such as measuring devices, hand gestures, writing systems, or iPhone applications – extra-neural resources which make us smarter by transforming the nature of otherwise difficult cognitive tasks into something more tractable for our biological brains. Cognitive artifacts allow us to distribute cognitive activities across time, space, and people (Hutchins, 1995). If the cognitively relevant interactions between brain, body, and environmental resources are sufficiently intimate and functionally

integrated in the performance of a task, it becomes increasingly arbitrary to single out the contributions of the body or the external world as mere “inputs” or “instruments” for cognition. In this case, or so the argument goes, cognition itself extends beyond the head into the world (Clark, 2008).

We consider the special case of group cognition rather than the general “extended mind” thesis for three reasons. First, humans are fundamentally social beings. We are deeply connected to each other, even before (developmentally and evolutionarily) we are connected to our tools. If humans lived in zoos, we would be classified as “obligatory gregarious” (Cacioppo & Patrick, 2008). We are rigged up to care about and depend upon one another. Specific hormones, such as oxytocin, have been implicated in trust, empathy and bonding (Domes, Heinrichs, Michel, Berger, & Herpertz, 2007; Kosfeld, Heinrichs, Zak, Fischbacher, & Fehr, 2005). Given the importance of people to people, groups of people are good places to look for emergent cognition. Philosophical discussion of emergent cognition has tended to focus on extensions to the individual human being via technological artifacts. Clark and Chalmers (1998) describe the parable of Otto, an individual who depends upon his notebook as an external memory aid. As closely connected as Otto and his notebook may be, the bonds of connectivity between people is far greater because we have adapted, over generations and within our lifetimes, to be sensitive to each others’ smell, sight, behaviors, creations, emotions, and thoughts. For this reason, we believe that groups may be one of the most promising places to look for extended minds that encompass individual people among their proper parts.

Second, people are becoming increasingly connected, and so determining whether and how collective intelligence is possible is becoming an increasingly important practical pursuit. The World Wide Web is a salient recent example of increasing connectivity. Mass-produced reference works such as Wikipedia make it clear that a large and decentralized collective has the capacity to organize information in very useful ways. Collective art projects such as Drawball,<sup>1</sup> SwarmSketch,<sup>2</sup> PixelFest<sup>3</sup> and HiveMind<sup>4</sup> demonstrate that groups comprising even 1000s of people can create unitary works of art that are at times coherent, representational, and take eventual forms that are not intended by any of the individual contributors. Although the technology underlying the Web is new, there is a deeper sense in which the Web is simply one of the most recent and socially significant manifestations of people’s perpetual drive to become more connected. Through innovations like the printing press, far-reaching transportation systems, and telecommunications networks, our lives have become increasingly intermeshed. In 1984, there were only about

1000 devices that could reach the global digital network. By 1992, about one million could. In 2010, over a billion can. From 1990 to 2003, mobile phone usage and global network usage each rose over 100-fold (Rehmeyer, 2007). While social critics have argued that people are less deeply enmeshed in their local communities than they used to be (Putnam, 2001), it is undeniable that people are becoming more broadly connected. Whether or not this has been good for individuals is debatable, but this question is mostly beside the point for the human collective. In fact, the speed and momentum of the collective’s push toward ever greater dependence suggests that it is a social force that is beyond the control of individuals to curtail. Richard Dawkins (1976) has argued that the interest of humans may be quite distinct from the interest of our genes, resulting in people behaving in ways that are opposed to their own good, but are for the good of their genes. Examples include people sacrificing their lives to save their kin, and people having kin in the first place even though they might believe it will reduce their happiness. Analogous to Dawkins’ selfish genes below the level of the individual, there are also “selfish teams” above the level of the individual. Both of these levels can cause individuals to behave against their own self-centered interests. Historically, the influence of selfish teams is precipitously increasing as societies become more organized and differentiated (Wright, 2001).

A third and final reason for focusing on group cognition is that it presents a useful case study for thinking about competition across explanatory levels. In most scientific fields, as a lower-level explanation gains traction, the higher-level explanation is often enfeebled or even literally undermined. Once basic genetics was well understood, explanations appealing to vitalism seemed unnecessary and unduly vague. Adams and Aizawa (2008) argue for a similar competition between levels of explanation between individuals and extended systems that contain individuals. They argue that in cases where it looks like an extended system (e.g., Otto plus his notebook) is cognizing, it turns out that only one part of the system (Otto) is doing the actual cognizing and the rest of the system is simply facilitating or enabling his cognition. This argument can be framed as a competition for explaining the system’s intelligent behavior. As Otto’s individual intelligence becomes increasingly plausible as a sufficient explanation for the larger system’s intelligent behavior, the need for positing a thinking system that includes Otto as only one part decreases in urgency. If a single part of a larger system suffices to explain the system’s essential behavior, then the very necessity of the larger system as an explanans is called into question.

Although competition between explanations at different hierarchical levels of a system frequently occurs, the relation between individual and group cognition is more complicated than this. In particular, for people within groups, there are situations where increasing the intelligence of the individual people can increase the coherence of the whole group and make it more likely to be an adaptive

<sup>1</sup> <http://www.drawball.com>.

<sup>2</sup> <http://swarmsketch.com>.

<sup>3</sup> <http://www.themaninblue.com/experiment/Pixelfest/>.

<sup>4</sup> <http://kevan.org/smaller.cgi>.

problem solving system in its own right. People frequently identify with the groups to which they belong, and important social group memberships become part of a person's self-identity (Smith, Seger, & Mackie, 2007). When this happens, people seek out ways to protect, perpetuate, and advance the group. When the people are motivated to promote their groups, then smarter people create more cohesive groups. Evidence from sports fans, political movements, families, and churches indicates that people are often so motivated (Haidt, Seder, & Kesebir, 2008).

A second possible positive interaction between the problem solving capacities of people and their groups is that smart people can modify their interactions to facilitate the group's interests. In particular, we create tools and protocols that allow us communicate with each other with increasing efficacy, thereby broadening the bandwidth of information exchange. As intelligent systems, individual humans seek out ways to create more coherent systems at higher levels than ourselves. We have instituted infrastructures such as peer-to-peer Internet protocols, chat rooms, Twitter, coffee houses, patent systems, professional organizations, and religious organizations to deepen the bonds that connect us to one another and promote exchange of information, beliefs, and innovations. Generally speaking, as a system increases the density of connections and dependencies among its parts, and decreases the density of connections and dependencies between its parts and parts outside of the system, it becomes an increasingly coherent entity in its own right, and more likely to be a practically indispensable unit of explanation (Gureckis & Goldstone, 2006). For example, a leading theory for the evolutionary origin of mitochondria and chloroplasts is that they were originally independent bacteria that became incorporated into the cytoplasm of cells, and once incorporated, conferred advantages for the cell because they allowed cellular respiration (mitochondria) and photosynthesis (chloroplasts) for energy production (Margulis, 1970). We now think of an entire cell, including mitochondria, as a coherent and explanatorily useful unit because of the strong dependencies between mitochondria and the rest of the cell. Just as mitochondria have been assimilated into cells, so people are assimilated into groups. Often, people even strive to self-assimilate themselves, creating indispensable group-level units.

Another complication that arises when parts make up a whole system, but the parts themselves are cognizant of the system, is that explicit cognition of the system by the parts can affect the parts' behavior. Typically, patterns of animal collective behavior emerge when individuals are interacting locally with one another, with no appreciation for the higher-order patterns that they are creating (Couzin, Krause, Franks, & Levin, 2007; Goldstone & Gureckis, 2009). Several human social structures do not fully fit this scenario, given that one self-professed motivation for individuals' behavior is to promote the group's welfare. For example, there has been a strong and growing movement to make software products, including the source code for

the software, available to any interested party without restrictions. A prominent reason often cited by programmers for contributing to this "open source community" is that they believe in the collective value gained by making software freely accessible and in the importance of the open source community itself (Lerner, Tirole, & Pathak, 2006). When the individuals that comprise a collective are capable of developing concepts of the collective, then the collective's identity and goal-directedness are intensified. The group formed by these individuals looks increasingly like a self-steering system.

A summary of these three kinds of interactions across person–group levels is that (1) people are typically motivated to have their groups flourish, (2) they alter their interactions to promote tighter within-group connections, and (3) they explicitly think about their group's welfare and how to perpetuate it. For these three reasons, an examination of person–group interactions reveals that there is not always a zero-sum competition between levels of organization, such that the more "unit-like" one level is, the less unit-like higher and lower levels are. In part because of the interactions across levels of organization, intelligent wholes can be associated with intelligent parts. Thus, we deny that increasing the cognitive capacities of individuals should make us less convinced that the whole group is capable of cognition. In fact, the above considerations make it likely that smart people create smart groups. This, however, begs the question of how we, individual humans, could spot a smart group of people in the first place, bearing in mind that a smart group of people is not necessarily the same thing as a group of smart people. In what follows, we lay out our proposal for how to identify cases of group cognition.

## 2.2. What's "emergent" about group cognition?

A considerable portion of the debate about extended cognition, including group cognition, has revolved around the concept of emergent properties at higher levels of organization: the potential for higher level structures to have properties that lower level components cannot have (Clark, 2008; Wilson, 2004). Thus, brains are assumed to have cognitive properties that no individual neuron has, or is capable of having, and this consideration is meant to suggest that groups have the potential to display emergent cognitive properties that no individual member has, or might even be capable of having. The implied psychological analogy between individuals and groups has a checkered history. Infamously known as the "group mind" thesis, the metaphorical use of mentalistic vocabulary to explain seemingly spontaneous forms of collective behavior (e.g., crowd movements, riots, fads, mass demonstrations) was a popular practice in the intellectual landscape of the late 19th and early 20th century (Le Bon, 1895/1960; for a historical discussion, see Wegner, Giuliano, & Hertel, 1985; Wilson, 2004, 2005). For many scientists and philosophers of that period, an appeal to "group minds" provided a suc-

cinct expression of what they perceived to be characteristic features of a group: their ability to function as a collective agent, the increased like-mindedness and conformity of people when they act as members of a group, and the *gestalt* of groups as emergent wholes that are “more than the sum of their parts.” To their own detriment, many traditional formulations of these ideas remained highly speculative, enmeshed with organicist metaphors, and often bordered on the occult. As a result, the “group mind” concept quickly fell out of favor with the rise of behaviorism in psychology, since it remained unclear where the “group mind” was supposed to reside, and how we could measure it (Wegner, 1986).

In the contemporary debate, the concept of *emergence* is frequently used – often in an ambiguous fashion – to denote one of three different (albeit related) facets of cognition at the group level (Theiner & O'Connor, 2010): its dependence on the organizational structure of a group, the manifestation of individually unintended cognitive effects at a group level, and the putative multiple realizability of cognitive properties by individuals and groups. Let us briefly illustrate each feature in turn.

We have said earlier that group cognition is not simply the unstructured aggregation of individual cognition, but the outcome of a division of labor among cognitive agents. The requirement of collaboration is necessary to avoid trivialization of appeals to group cognition. Consider a dyadic group that is composed of two individuals – Angelina and Brad. Angelina can expertly solve a Rubik's cube, whereas Brad cannot even solve a Rush Hour puzzle. If we countenance the idea of group cognition, do not we have to say that the dyad as a whole – the “Brangelina” group – can automatically solve a Rubik's cube courtesy of Angelina's cognitive abilities? This may amount to a limiting sense in which groups can be said to cognize, but only in a strictly derivative (and thus dispensable) sense insofar as the cognitive abilities of the group are entirely parasitic on the cognitive abilities of its members. Surely the interesting cases of group cognition – including the ones on which we shall focus – are quite unlike that, insofar as their manifestation is an “emergent” effect of the organizational structure adopted by the group. The relevant sense of *emergence*<sub>1</sub> here can be defined as a failure of “aggregativity” in the sense of Wimsatt (1986). 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. Emergent<sub>1</sub> properties which fail to satisfy all four of these conditions are only “minimally decomposable” in the sense of Simon (1969). In contrast, the purely aggregative ability of “Brangelina” to solve a Rubik's cube can be exhaustively explained in terms of Angelina's ability to solve a Rubik's cube.

A different sense of *emergence*<sub>2</sub> reflects the puzzlement that we experience when we observe apparent conflicts between the properties of entities at different levels of organization. For instance, how can a traffic jam move backwards when all of the cars comprising the jam are moving forward or are stopped (Wilensky, 1997)? How can infusing heat to the movement of iron molecules, adding randomness in gradually reducing increments, cause the molecules to achieve a more orderly crystalline arrangement than if the heat had never been supplied? In such cases, puzzles arise because similar properties are manifest at multiple levels, albeit with opposite values. Both cars and traffic jams have movement directions, although they may move in opposite directions. Both individual iron molecules and entire configurations of molecules can be described in terms of their orderliness, even though less order at the molecular level may result in more order at the crystalline level.

The mismatch between properties of systems and their parts is all the more puzzling when it is we, individual humans, who are the elements, and when the system that contains us is a group of people embedded in an environment. We marvel at situations in which public benefits flow from selfish intentions (Smith, 1776), or in which public vices unexpectedly spring from private virtues (Hardin, 1968), because those situations violate our naïve intuitions about how local behavioral rules scale to the global properties of inter-connected wholes (Resnick, 1994). More generally, we speak of emergent<sub>2</sub> cognitive properties of groups which arise from the local interactions between many individuals, but without being planned or purposefully designed by any of these individuals (or some central planning agency), and which those individuals may even fail to notice.

The third and final sense of *emergence*<sub>3</sub> refers to our contention that groups and individuals are both aptly viewed as units of flexible cognizing, because they share important informational processing commonalities. In the philosophy of mind, the thesis that functionally equivalent cognitive states and processes can be realized by distinct physical mechanisms is commonly referred to as *multiple realizability* (Block, 1997; Fodor 1974). Critics of extended cognition have argued that describing supra-organismic entities in terms of emergent<sub>3</sub> cognitive processes such as memory, problem solving, and attention risks making these notions too nebulous and that the expanded construals no longer pick out natural kinds that allow fertile inductions (Adams & Aizawa, 2008; Rupert, 2004). In contrast, we are impressed by the extent to which units that are composed of different kinds of stuff are organized in functionally similar ways (see also Clark, 2008; Hutchins, 1995). The inductive power of grouping together superficially dissimilar systems should not be too surprising given the large literature on systems composed out of very different materials showing the same patterns of organization. For example, in many systems, individual elements move randomly. If a moving element touches another element, it becomes

attached. The emergent result is a fractally connected branching aggregate. This process, called diffusion-limited aggregation (DLA), has been implicated in the growth of human lungs, frost on glass, and cities (Ball, 1999). Creating a category of DLA phenomena is scientifically useful because the resultant overall patterns of thin, fractally connected branches, are remarkably similar, and have almost identical statistical properties. Once frost is appreciated as an example of DLA, one can predict how it will change with time, how it will be affected by temperature, what shapes it can and cannot attain, and so on. Physicists often use the term *universality* to describe generic phenomena (such as DLA) which happen in the same way across a wide range of physically diverse systems.

Because the functionalism that we endorse here is of a limited kind – i.e., we do not assert that individual organisms and groups instantiate *identical* cognitive properties – we are not committed to any strong theses about multiple realizability. One may, of course, describe cognitive traits very abstractly – in terms of memory, reasoning, etc. – and claim that both individuals and groups instantiate those (same) traits, but such characterizations are likely to be relatively unproductive for prediction or explanation (Allen, 2002). We are sympathetic to the point (e.g., Bechtel & Mundale, 1999) that philosophical arguments for multiple realizability only appear to work because they ignore the different grain of analysis between higher-level properties and lower-level realizers. Nevertheless, abstract, functional characterizations of the phenomena that elide certain details provide useful organizing principles for the investigation of cognition. By this we do not mean to rule out the possibility that the very same cognitive property could be possessed by an individual and a group, but it is vanishingly unlikely that a group-level cognitive property would be functionally indistinguishable from an individual cognitive property.

Many information processing capacities, including memory, attention, and problem solving are good candidates for being inductively powerful categories even when expanded beyond systems composed of neurons precisely *because* they concern information. What matters for information is pattern and function rather than specific material instantiation. A source, like a bootleg recording of The Clash's first gig, opening for the Sex Pistols in 1976, retains much of its precious information no matter what physical form it takes – analog magnetic polarizations on a cassette tape, grooves on a vinyl LP record, optically reflective pits on a DVD, or electrically encoded digital 0s and 1s in a MP3 file – although they are far from being identical reproductions of the original. No less than natural patterns like DLA, positive feedback, negative feedback, and stigmergy, information processing patterns recur across many materials because of their stability and utility. Of course, the proof is in the pudding here. As proponents of group cognition, we will discuss in Section 3 some observed cases of group-level information processing that we take to reflect important properties of cognitive systems. We will stop

short of claiming that the cases of group behavior unambiguously demonstrate group cognition. Instead, our claim is that the groups show properties critical to cognitive systems that are different from those shown by the individuals within the groups.

### 2.3. GLAD to cognize

As our above survey indicates, it is not surprising that the concept of emergence has figured prominently in recent debates over socially or environmentally extended cognitive systems. To our knowledge, however, no one in this debate has considered the inverse notion of “demergence:” the situation where larger scale structures not only lack the properties that their components have, but are somehow constitutionally incapable of having those properties. On the traditional view of consciousness, it is a demergent property – one that pops into existence at a certain level of organizational structure, but then disappears from higher levels of organization except insofar as it is found in the parts.

In the previous section, we gave the example of the direction of motion of a traffic jam as an emergent property that is puzzling because its direction can be the opposite of all the parts – the cars making up the jam. This, of course, is a general property of waves – the molecules of water through which a wave travels may actually be traveling in the opposite direction to the wave, a fact exploited by anyone who has learned how to paddle a surfboard away from the beach. But direction of motion is not a demergent property: once it is found at a given level of organization, it is also found at higher levels, even if its value need not be a simple average or aggregation of the value of the parts. If emergent properties are puzzling (and we remain neutral on whether it is an ontological puzzle or merely an epistemological one), we suggest that demergent properties should be even more so. If emergence depends on organization, how could the addition of more stuff block the same kind of organization at larger scales? Our tentative suggestion is that a default answer to this question should be that it cannot. That is, once the resources exist to generate a new phenomenon at some scale, larger collections should be assumed to have the same potential. Thus, if cognition depends on the organizational structure within individual organisms, groups of organisms should be assumed to have the same kind of capacity for organizing themselves, and thus the same potential for instantiating cognitive properties. Let us call this assumption the principle of *group-level anti-demergentism* (GLAD).

From an epistemological point of view, GLAD has important implications about what burden of proof must be met by those who deny that group cognition is a real world phenomenon. To be sure, the critics we are primarily concerned with in this piece are not strict demergentists – they will readily concede that cognitive properties could in principle be possessed by super-organismic entities, yet they will not admit that there are any actual instances to be found. It seems to us that these critics are demanding

more than the empirical success of applying the methods and models of cognitive science to groups (or other extended entities), which we shall illustrate further below. For those critics, who think that cognitive properties are at least *de facto* demergent, GLAD equally implies that they need to do more than press their philosophical intuitions about the ontological privilege of individuals being the sole bearers of cognitive properties. They owe us an explanation of *why* cognitive properties do not behave like many other properties that can be instantiated at multiple levels in the organizational hierarchy.

By proposing the GLAD principle, we thus hope to re-orient the current debate about parity considerations of the sort originally introduced by Clark and Chalmers (1998) in their philosophical defense of the “extended mind” thesis. In support of their claim that cognition extends beyond the head into the world, Clark and Chalmers argued that if, in confronting some task, a part of the world functions as a process which, were it done in the head, would be accepted as part of a cognitive process, then that part of the world *is* equally part of the cognitive process (ibid., p. 8). In earlier work, one of the present authors proposed an analogical extension of this inference by way of a “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 (Theiner, 2008; see also 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 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.

Highlighting the first role, Clark (2005, p. 2, fn. 3) has instructively compared parity to a “veil of metabolic ignorance,” similar to the role which the “veil of ignorance” famously plays in Rawls’ (2001) social-contract theory of justice. The point of Rawls’ construction is to establish an impartial position from where we can determine what principles of justice people would jointly accept if they were negotiating as free and equal agents behind a veil of ignorance, i.e., deprived of any information about their personal characteristics (e.g., their ethnicity, gender) and their particular social and historical circumstances. In order to perform this function, the veil of ignorance itself need not – in fact, should not, so as to avoid circularity – embody a substantive conception of justice. In our argument, GLAD acts as a “veil of ignorance about scale”

which limits the kind of information to which one can justifiably appeal in assessing the cognitive *potential* of groups. (More precisely, we think that GLAD suffices to establish the *nomological possibility* of group cognition as a metaphysical default, although we will not argue this point here.) But without an independent account of cognition, GLAD cannot be used to infer that groups *actually* have cognitive properties, nor even to assess the likelihood that they do. Neither can it be used to determine what those cognitive properties would be.

Let us briefly illustrate the burden-shifting role of GLAD by responding to a challenge for proponents of group cognition raised by Wilson (2004). Concerning the aforementioned analogy that individuals are to groups like neurons are to the brain, Wilson (ibid., p. 291f) points out that “individual neurons do not *perceive* (but fire in response to a stimulus), *remember* (but transmit information about the past), or *plan*. By contrast, individual agents do all of these things, and these are just the sorts of properties attributed in socially distributed cognition.” Wilson’s point is that we do not attribute the same kinds of properties at the neural and agent levels, so there is something inherently suspicious about attributing similar kinds of properties at both the individual and group levels. We agree that it would be a compositional fallacy to infer that groups have cognitive properties simply on the basis that their members have cognitive properties, just as it would be a mistake to infer that the finance committee is male just because all of its members are male. However, because of our commitment to GLAD, we also think that just pointing out that individuals have cognitive abilities should not predispose one against the claim that groups also have cognitive abilities.

In this context, it is important to stress that GLAD should be taken to refer to generic *types* of properties, not to specific property *instances*. As with the example of direction of movement, there’s no guarantee that the actual values of the cognitive properties of the larger groups will be the same as those of the individuals composing the group, nor do we accept that any property that is actually attributed to an individual automatically is a property of higher level structures to which that individual belongs. (As a case in point, since motion is always relative, the direction of motion is in fact demergent at the scale of the entire universe.) So while we can talk about groups solving problems, for example, there is no guarantee that groups can solve the same problems as the individuals that make them up, or that if they can solve the same problem that they will come up with the same solution that an individual would.

Edwards and Pratt (2009) describe a striking case of this dissociation between individual and group problem solving in their study of ant colonies. A hallmark principle of neo-classical theories of rationality (known as *regularity*) is that a decision-maker who prefers option A for its intrinsic benefits over option B should not switch this preference based on the presence of a third “decoy” C. It is well known, however, that human beings are prone to regularity violations in the presence of third options which are dominated



by one original option, but not the other. Edwards and Pratt show that unlike humans, ant colonies are immune to the “decoy effect” when they have to make collective choices between different nest sites. Notwithstanding this difference in performance, such a comparison logically requires that we consider both individuals and groups alike as subjects of decision-making.

This finally brings us to the question of which cognitive or mental properties are in fact manifested by groups of cognitive agents. Answering this question requires attention to the second epistemic role that we have distinguished, which is to supply criteria for the attribution of particular cognitive or mental properties. Like the ancestral parity principle from which it is derived, 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.

Instead, 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. 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. In this paper, we focus on group problem solving, following the research of Goldstone and colleagues, and on group memory, showcasing research on transactive memory systems. Given that the problem solving capacities and memory dynamics of groups are not identical to those of individuals, we believe we can sidestep the issue of cross-level explanatory competition due to Adams and Aizawa (2008). Being a member of a group may indeed facilitate individual cognition. Nevertheless, the group as a whole exhibits cognitive capacities that are different from the aggregated cognitive capacities of individuals, and if the group's cognitive capacities can be recast as enhancement of cognition in the individuals belonging to the group, then this is a post hoc redescription that fails to respect the actual causal structure of the capacity.

### 3. Case studies of group cognition

#### 3.1. Stigmergic path formation

Groups often create path systems that are mutually advantageous to the members of the group without that

being the intent of any of the members. Ants create efficient foraging trails and people make efficient trails interconnecting university buildings. In both life forms, these trail systems arise because agents are motivated to take advantage of the trails left by their predecessors. In the process of exploiting previously left trails, agents further reinforce these trails, potentially leading to a lock-in of originally tentative and faint paths. Early trail blazers through a jungle use machetes to make slow progress in building paths – progress that is capitalized on and extended by later trekkers, who may then widen the trail, then later put stones down, then gravel, and then asphalt.

The “Active Walker” computational model accounts for this dynamic by assuming that walkers' movements are a compromise between going to their destinations and going where the travel is easiest due to previous walkers' steps. (Helbing, Keltsch, & Molnár, 1997; Helbing, Schweitzer, Keltsch, & Molnár, 1997; Moussaid, Garnier, Theraulaz, & Helbing, 2009). Predictions of the model have also been confirmed by laboratory experiments with humans (Goldstone & Roberts, 2006). Furthermore, parametric variation within the Active Walker model can potentially be used to guide policy decisions. For example, two of the most important parameters of the model are the visibility of paths (the extent to which an agent is influenced by distant patches' travel ease) and a path's decay rate (how quickly the influence of a step on a patch's ease of travel dissipates). If path decay is set to a high level, then the paths that agents make disappear relatively rapidly. The Active Walker model suggests specific interventions depending upon a community's goals. For situations where conserving the total amount of pathway is desirable (e.g., when valuable vegetation is destroyed by paths), the model's advice is that planners should explore ways of increasing either path visibility or path decay (Goldstone, Jones, & Roberts, 2006). While increasing visibility intuitively makes a group better coordinate itself to create and exploit good paths, it may be less intuitive that making paths decay relatively quickly promotes better path systems. However, without decay, a group's path system will become stuck on the early paths created by walkers, and these paths will be bee-line paths connecting destinations rather than path networks that efficiently combine trails that are close. As shown in Fig. 1, the most efficient path network, shown in 1A, is achieved by a self-organized collective with highly visible, but short-lasting paths. One moral for abstract paths (i.e. innovations) is that intentionally building in obsolescence into paths (and speculatively, computer architectures, patents, and music) may promote social structures that build upon earlier innovations, but are not unduly locked into their specifics.

The important point is that the computational model and group experiments suggest, and even converge upon, ways for a group as a whole to create effective path systems. These path systems are solutions to a problem for the group, not the individuals. The individuals, at least those specified by the model, are simply taking steps that




Fig. 1. The influence of parameters on path formation within the Active Walker model (Helbing, Keltsch, et al., 1997), as reported by Goldstone et al. (2006). The destinations are shown by colored circles. The darkness of a patch is positively related to its level of comfort for walking, and indicate the eventual paths after 1000 iterations. (A) Decay rate of a path = .1, impact of each step on comfort level = 1000, path visibility = 100, (B) decay rate = .1, impact = 1000, visibility = 1, (C) decay rate = .001, impact = 10, visibility = 10, and (D) decay rate = .1, impact = 1000, visibility = 10. The most efficient path system is found in Panel A, in which walkers are strongly influenced by the comfort level of the patches (visibility of paths is high), but the influence of steps on travel comfort quickly dissipates (paths decay quickly). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

compromise between going where they want to go and where others have gone before. The group's problem is different – creating an efficient path system for all members. Furthermore, the role of a strategic human planner in establishing the group's parameters can be removed. If groups naturally differ in the built-in obsolescence of their paths, then they will also differ in the efficiency of their path systems by the above analysis. Groups that stumble upon path obsolescence will create better path systems and might be expected to spread their structure by either evolutionary or cultural means. 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 (in fact, it is NP-hard), 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. Furthermore, there is excellent reason to believe that this problem solving approach is found in many groups because it is based on a common pattern of stigmergy – indirect communication between agents that is achieved by agents modifying their environment and also responding to these modifications (Dorigo, Bonabeau, & Theraulaz, 2000; Marsh & Onof,

2008). It is employed by Amazon.com automated book-suggesting bots, ants, robotic builders, technology companies, and people making trails in the snow. 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.

### 3.2. Collective Coloring Problem

Another good example of groups solving a difficult collective problem is Kearns, Suri, and Montfort's (2006) network "coloring" problem in which a group of people are arranged in different topologies of social networks. The participants' collective task is to choose their individual colors such that no network neighbors share the same color. Two network topologies, "cycle" and "preferential attachment" networks, are shown in Fig. 2. The groups were typically able to find admissible color patterns, with some topologies, such as the preferential attachment network, being more difficult than others, such as the simple cycle. On some problems, participants could see the dynamically changing colors of the entire network. On other problems their visibility was limited to their immediate network neighbors' colors. One surprising result was that for the difficult preferential attachment topology, faster solutions were found when visibility was limited to neighbors, rather than being network-wide. The authors suggest that participants in the full visibility condition may have suffered from information overload due to the complexity of the preferential attachment network.

For our present purposes, this experiment makes two important points. First, 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




Fig. 2. Examples of network topologies used by Kearns, Suri, and Montfort's (2006). In the Cycle arrangement, a circle of group members is established, and members are each connected to their two neighbors. A possible solution is shown in which each member who has colored themselves black is next to two neighbors who have colored themselves white, and vice versa. In the preferential attachment configuration, a network is created by adding nodes that tend to link to existing nodes that are already highly connected. This configuration requires at least three colors, and a three-color solution is shown in which no neighboring members share the same color.

interpreted as requiring relatively sophisticated problem solving strategies including planning, back-tracking, and multiple constraint satisfaction. 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. 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. Mason, Jones, and Goldstone (2008) describe another situation in which providing individuals with more information leads to worse group-level performance.

### 3.3. Acquired division of labor within groups

In another collective behavior paradigm, Roberts and Goldstone (2009; submitted for publication) examined the mechanisms by which the members of a group coordinate their activities to reach a specific group outcome. Their first desideratum was to have a situation in which there is a pure group goal, and everybody is given that same goal. Given that they were especially interested in the dynamics of coordination, a second desideratum was for there to be multiple rounds over which coordination could occur, as well as multiple replications of the task. Finally, they wanted to create a situation in which the group benefits from different people learning to play different roles within the group, in an acquired division of cognitive labor. The game they developed is called "Group Binary Search." One hundred and six participants were divided into 18 Internet-connected groups, and the members of each group were told that a computer would be randomly choosing a target number between 51 and 100. During each 15-s round, each participant entered a number between 0 and 50, and the computer summed these numbers and compared this sum to its target number. In the "numeric feedback" condition, participants were told how high or low their group's sum was. For example, if Wolfgang, Ludwig, and Johann were in a group and typed in 25, 17, and 12 respectively, and the computer's target number was 72, then each member of the group would be told that their sum was too low by 18. In the directional feedback condition, participants were only told whether their group's sum was too high or too low. After each round, participants got the chance to react to the computer's feedback and alter their individual numbers. This guess-then-feedback process was repeated for as many rounds as was needed for the group to reach the target number. Five games each of directional and numeric feedback were conducted with each group, making for an entire experiment session that lasted about 50 min.

Group performance generally increased as the number of members in a group decreased. With only one participant, the problem degenerates to a simple "binary search" task with at most  $\log_2(N)$  rounds needed to find the target number among  $N$  numbers in the directional feedback case and only two rounds needed with numerical feedback. As shown in Fig. 3, groups improved over the course of their

sessions. One of the reasons for this improvement is that group members become strategically differentiated over rounds within a game, and across games. The degree of strategic differentiation is well captured by measuring participants' *reactivities* – how much they change their numbers in response to the computer's feedback. For example, if the group was told that their sum was too high, and a member subsequently lowered her proffered number from 37 to 32, then her reactivity would be coded as +5. Using this measure, we obtained two pieces of evidence supporting the spontaneous differentiation of people into different strategies. First, for large groups, over rounds of games the variability of reactivities *within* a participant decreases. A participant tends to either become the type of person that reacts, or does not react, to the computer-supplied feedback. That is, participants become increasingly consistent, and hence predictable by others. Second, for large groups, the variability of reactivities *across* people increases over games. This indicates that some members are adopting roles of reactors, and others of non-reactors. Furthermore, the extent to which reactivities become more diverse across people is positively correlated with a group's performance. As people become more diverse, the group does better. In sum, the group comes to be able to solve a difficult coordination problem because the variability within individuals' behavior is decreased, and across individuals is increased.

The Group Binary Search is a good example of the members of groups adapting to become more differentiated, thereby increasing the coherence and effectiveness of the group. Computers, cell phones, and PDAs do adapt to us, but not nearly as effectively as we adapt to each other, because we are smarter and we are built to care about each other and fitting in. Interestingly, the influence of group members on each other in this experiment leads to differentiation. Most social psychology experiments have focused on the complementary group influence – whereby innovations or norms pass from member to member via




Fig. 3. Over the course of five successive games, groups improve when given either numeric or directional feedback, reaching the target sum in fewer rounds with time. Groups also perform better with the more informative numeric feedback compared with the directional feedback.

imitation, leading to greater similarity and assimilation of the group members (Mason et al., 2008). The experiment is too simple to be an analog of the kind of division of labor that is found in science, businesses, and governments. However, it provides an idealized dynamic for how members of a culture or community become differentiated. For the present purposes, 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.

### 3.4. Remembering in groups

When people remember in groups, do they ever form a collective memory system in its own right? Harking back to the traditional notion of a “group mind,” the attribution of memory to groups has at times been justified by the holistic slogan that the whole is more than the sum of its parts (Wegner, 1986). In social psychology, the potential of groups to integrate the contributions of their members in a manner that produces outcomes of a higher quality than would have been achieved by simply adding up individual members' efforts is known as the *assembly bonus effect* (Collins & Guetzkow, 1964) or as *process gains* (Steiner, 1972). Curbing our enthusiasm, Pavitt (2003) has recently argued that groups tend to remember less than the sum of their parts, and this gives us reason to be “reductionists” about group memory. Groups remember more than any one of their members, to be sure, but these performance gains are purely aggregative – they do not stem from interactions. If we compare the performance of real groups with baseline models that add the non-redundant amount of information recalled by individuals working alone (Lorge & Solomon, 1955), real groups consistently underperform nominal groups. Pavitt's critique has stirred an interesting controversy which provides a springboard for our discussion.

There is indeed robust evidence of *collaborative inhibition* (Weldon & Bellinger, 1997) in collective recall tasks. Clearly, this shows that opportunities for group interaction alone do not guarantee optimal group performance. A frequently cited explanation for this effect is that listening to the recollections of others disrupts one's own retrieval strategies (“production blocking;” Basden, Basden, & Henry, 2000; Diehl & Stroebe, 1991). Other factors considered to hamper group performance include social loafing and coordination losses (the “Ringelmann effect”), com-

petitive motives of individuals for withholding potentially relevant information, evaluation apprehension, a tendency to rehearse shared information as a means of receiving social validation, and a failure to develop a differentiation of expertise within a group (Hollingshead & Brandon, 2003; Wittenbaum, 2003).

However, a less than optimal recall performance need not undermine the explanatory value of contemporary theories of memory as a group-level phenomenon. For one, process losses are no less significant than process gains to underscore the constitutive nature of social interactions and communication processes for group memory. To study collective remembering as a group process, we do not have to assume that communication is a flawless conduit for pooling individuals' memories (Propp, 1999; Wegner, 1986). Moreover, Pavitt's argument, which is characteristic of individualistic approaches to group memory, is based entirely on a simple input–output model of group performance. Such a viewpoint has important drawbacks. First, if we only look at aggregated performance data, process gains and process losses may cancel each other out. Second, the comparison with individualistic baseline models, while useful in certain respects, is implicitly biased towards simple tasks in which teamwork is strictly optional. This differs from the complex cognitive requirements for memories which in real-world situations are always processed and maintained, in a distributed fashion, by groups, institutions, or even entire cultures (Argote, 1999; Donald, 1991; Walsh & Ungson, 1991). Third, the narrow focus on intellectual tasks promotes a reliance on performance measures (e.g., *amount recalled*) which may not be suited to capture other integrative effects of collaborative memory processes in groups (Propp, 2003).

To study the functional organization of memory in couples, families, and small workgroups, Wegner and colleagues (Wegner et al., 1985; see also Wegner, 1986, 1995) introduced the notion of a *transactive memory system* (TMS). Their main interest was to analyze “processes that occur when the TMS is called upon to perform some function for the group – a function that the individual memory system might reasonably be called upon to perform for the person” (ibid., p. 256). We can distinguish two components of a TMS. Its *representational* component consists of the organized stock of memories that are retained by individual members, including higher-order memories about who knows what. The creation and maintenance of these representations is performed by the *procedural* component, which includes all direct and indirect communication processes (“transactions”) by which individuals cooperatively allocate, encode, retrieve, share, and elaborate memories. For instance, effectively managing a TMS requires regular updates of the transactive memory base, allocating incoming information as well as encoding responsibilities to the relevant experts, and coordinating the retrieval of unshared individual memories. For a group to develop successful transactive memory procedures, it is imperative that its members learn who to ask, how to ask, and how to cue

each other appropriately. Understanding the patterns of transactions which occur when people remember together as a group, and under which conditions they lead to collaborative inhibition or facilitation, is thus a distinctive goal of process-oriented approaches to group memory (Theiner, 2010).

For instance, Wegner, Erber, and Raymond (1991) found that long-time couples learned and recalled more words than strangers, because they were able to divide their encoding responsibilities according to their tacit knowledge of each other's expertise. Extending their study, Hollingshead (1998a) showed that strangers who could communicate during encoding later recalled more collectively than when they could not. On the other hand, dating couples performed worse in the same condition, because communication disrupted their pre-existing TMS. Hollingshead (1998b) also showed that dyads who were allowed to communicate during both learning and retrieval suffered no significant process losses compared to nominal dyads, which suggests that dyads recall best when they can collaboratively learn their own group-level strategies. Interacting dyads of friends suffered fewer process losses relative to nominal dyads than did strangers (Andersson & Roennberg, 1996), and long-term married couples with a thoroughly established TMS showed no signs of process losses at all (Johnansson, Andersson, & Roennberg, 2000). In a recent study of autobiographical memories in elderly married couples, Harris, Keil, Sutton, and Barnier (2010) found that the use of shared retrieval strategies, interactive cuing styles, and repetition promoted collaborative facilitation, whereas unevenly distributed expertise, strategy disagreement, and recall of extraneous information were associated with collaborative inhibition.

Wegner's conception of a TMS has been used to analyze the productivity of small workgroups performing collaborative assembly tasks (Moreland, 1999; Moreland, Swanenburg, Flagg, & Fetterman, 2010). In this line of research, TMS is treated as a latent group-level variable that is manifested in terms of *memory differentiation*, i.e. the tendency of group members to specialize in recalling distinct aspects of the assembly process; *task credibility*, i.e. how much members trusted one another's expertise; and *task coordination*, i.e. the effective use of transactive retrieval strategies (see also Lewis, 2003). Liang, Moreland, and Argote (1995) showed that work groups performed better when their members were trained on the task together rather than individually, and these performance gains were mediated by the impact of a TMS. Members who were trained together also acquired more elaborate, accurate, and mutually shared beliefs about the distribution of know-how in the group. Examining how groups adapt to changing task demands, Lewis, Lange, and Gillis (2005) found that groups with an intact TMS and experience with at least two tasks in the same domain showed higher levels of learning transfer than groups with no prior TMS.

In sum, our considerations suggest that Pavitt's argument against group memory is inconclusive. The lack of

evidence for a consistent "assembly bonus effect" is not a reason to abandon the study of collective remembering as a group-level process. Instead, it helps us to better conceptualize and measure the role of communicative interactions for the organization of memory in groups. 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.

### 3.5. Conclusions to case studies

Integrating these four case studies offers some generalizations about the kind of flexible problem solving that is surely one of the core components of thinking systems, be they composed of neurons or of people (each equipped with a brain composed of neurons). Different group-level problems require different interaction patterns. At first sight, the stigmergic dynamic at work with path formation and innovation propagation seems at odds with the differentiating dynamic underlying a transactive memory system, the Collective Coloring Problem, and Group Binary Search. Underlying the bona fide surface difference between people behaving in increasingly similar versus dissimilar manners, there is a deeper commonality uniting these case studies. They all show mutual adaptation of people to each other, with group-level consequences that can be placed under selection pressures. The simultaneous and mutual influence of people on each other makes it impossible to simply treat one person as the "prime cognizer," and the others as mere cognition supporters. As the extent of interactions among people increases, so does the impetus to ascribe functionality to the entire assembly. The fact that some group-level outcomes are better than others, and the existence of mechanisms for groups to perpetuate and extend upon these better outcomes, provides grounds for believing that the group will become, over time, an increasingly indispensable level of description.

A final issue that is worth revisiting with the specific case studies in our mind is: "Why do these group-level phenomena require an explanation that goes beyond individual cognition?" A first response is to emphasize that the groups' behaviors depend critically on environmentally-enabled interactions between people. For example, the different pathways observed in Fig. 1 vary as a function of environmental differences in the ease of creating paths and the rate of path decay. Variables such as these are part

of the larger context in which people are situated, and are not reducible to individual cognition. Second, the groups' solutions to problems may never be entertained by any individual. Fig. 1A shows a good solution to the problem of finding a short path system that connects the destination points. However, this group-level solution that may not be considered by any individual, and in fact, the individuals may not even know that this is a problem they should be trying to solve. The individuals are typically making short-sighted and selfish choices about which direction of travel will be least costly. Third, the groups' solutions are not simply composites, unions, or intersections of individuals' solutions. Fig. 1A is a global solution to a "find the shortest path set" problem, but it is not put together by combining whole solutions or parts of solutions from individuals. In typical brainstorming sessions, group solutions surface by cobbling together individually entertained and offered solutions. In contrast, for at least the path formation and collective coloring case studies, global solutions are only available at the group level, not the individual-level.

One of the clearest reasons why the groups' behaviors in the case studies are not reducible to individual cognition is that the groups' adaptability and information processing capacities arise from interactions among the individuals and their environments. Exactly the same individuals, configured in different network topologies (Mason et al., 2008), or placed in an environment with different responsiveness (Goldstone & Roberts, 2006), are involved in qualitatively different group patterns. Furthermore, these group-level patterns can be aptly evaluated from problem solving and information processing perspectives. For example, it makes sense to ask about a pattern whether the group is producing a good or optimal solution, whether it is located in a local but not global maximum, and what kind of representations and internal states are implicated by the pattern.

#### 4. Resolving the "coupling-constitution" issue

A major point of contention in the current debate over the "extended mind" thesis is the question of the conditions under which the causal intercourse between brain, body, and environment is intimate enough to spawn the existence of extended cognitive systems. According to Adams and Aizawa, the "coupling-constitution fallacy" is "the most common mistake extended mind theorists make" (2010, p. 68) when they tacitly move "from the observation that a process  $X$  is in some way causally connected (coupled) to a cognitive process  $Y$  to the conclusion that  $X$  is part of the cognitive process  $Y$ " (ibid., 2009, p. 81; see also Adams & Aizawa, 2008, Chapters 6–7). As a non-cognitive example where said fallacy should be apparent, they point out that the expansion and contraction of the bimetallic strip in a thermostat is closely coupled to the ambient temperature of a room, yet this gives us no reason to consider the expanding and contracting strip as extending into the

room (Adams & Aizawa, 2009, loc. cit.). A similar criticism is made by Rupert (2004, pp. 395–396), who argues that even if one grants that we cannot fully comprehend many episodes of human cognition unless we consider the wider context in which they occur, it does not follow that the embedding context is therefore part of cognition. For instance, to understand the historical dimension of Nazi Germany's invasion of Poland, one has to know a great deal about the economic conditions in Germany during the 1920s, but this does not make the latter (mereological) parts of the invasion.

We do not think that we are guilty of committing this sort of fallacy. To begin with, when we claim that an individual cognitive system  $X$  is in some principled way coupled to another individual cognitive system  $Y$ , we do not mean to imply that  $X$  is thus part of  $Y$ . Instead, what we assert is rather that the individuals who instantiate  $X$  and  $Y$  can engage in structured interactions so as to constitute an organized group-cognitive system  $Z$  that encompasses those individuals among its proper parts. Hence none of the above examples is particularly apt to bring out what we take to be the critical issue: when do we have good reasons to posit such a system  $Z$ , and what explanatory role does an analysis which refers to properties of  $Z$  play? Since we can hardly ascertain the epistemic value of such explanations – and the systemic taxonomies which they presuppose – on apriori grounds, we agree with Rupert (ibid.) that it would be counterproductive to endorse an unrestricted version of this principle. This would quickly lead either to a gratuitous proliferation of mostly ephemeral cognitive systems, or the monist conclusion that there is really only one huge cognitive "mind-meld," assuming that everything is somehow contextually dependent on everything else. Although we should not outright dismiss either possibility on the spot, the onus here is on the extended/group mind theorist to tie the occurrence of group cognition to specific and sufficiently robust forms of collective behavior.

For a group of two or more people to constitute a cognitive system in its own right, we require that these people are coupled (in their functioning as members of the group, collectively performing a cognitive task) so as to form an *integrated system* with *functional gains*. Following Wilson (in press, p. 19), we can break down the composite notion of a *functionally gainful, integrative coupling* as follows. First, two (or more) elements are *coupled* just in case they exchange information by means of reliable, two-way causal connections between them. Individuals who are collectively coupled are interdependent in their cognitive and behavioral activities. Second, two (or more) coupled elements form an *integrated system* in situations in which they operate as a single causal whole within the causal nexus – with causes affecting the resultant system as a whole, and the activities of that system as a whole producing certain effects. Third, an integratively coupled system shows *functional gain* just when it either enhances the existing functions of its coupled elements, or manifests novel

functions relative to those possessed by any of its elements. Distinguishing these two aspects of functional gain is important for an account of group cognition, because it implies that the cognitive interdependence between people has both individual-level as well as group-level effects. In our case studies, we have carefully described a variety of situations in which groups function as integratively coupled cognitive systems in this sense. We have shown that first, groups exhibit emergent (sensu Wimsatt, 1986) cognitive capacities that are different from the aggregated cognitive capacities of its members; second, individuals act and cognize differently as members of the group than they would in isolation.

In order to forestall a potential misunderstanding of our account, let us reiterate that we are not aiming to state a “mark of the cognitive” here. We find ourselves in agreement with Adams and Aizawa (2009, p. 75) that the logic of the “coupling-constitution” argument is an orthogonal issue. Following our earlier methodological recommendation, the term “cognitive in the previous paragraph does not refer to an underlying abstract notion of cognition, but merely acts as a placeholder for particular cognitive capacities. The slight loss of generality is outweighed by a notable gain in substance if we dispense with the word “cognitive,” and replace it with more specific terms such as “problem solving” or “memory.” Consequently, we see no harm in our parenthetical reference to the tasks for which those capacities are harnessed as “cognitive tasks.” While there are certainly important family resemblances among those capacities which allow us to group them together under a single generic umbrella, we do not have to assume that there must be a “hidden essence” of the cognitive for such a categorization to be useful. On the other hand, our discussion suggests that specific information processing capacities such as memory, attention, problem solving, and learning are indeed inductively powerful categories for analyzing distinctive patterns of both individual and group behavior. Similarly, since we do not intend to provide a “mark of the social,” our use of the terms “group” and “collectively” merely reflects the explanatory role which those terms play in the models we have discussed. If anything, we have tried to give a “mark of a [cognitive] system,” by pulling together several conditions under which it is profitable for cognitive scientists to study groups as adaptive problem solving units or as transactive memory systems in their own right.

But the dispute over group cognition does not end here. Perhaps we are guilty of committing an equivalent of the “systemic” version of the coupling-constitution fallacy (Adams & Aizawa, 2008, Chapter 7), which differs from the simple version by its extra appeal to the notion of a cognitive system. The first step of the criticized argument moves from the premise that there is an integrative coupling between brain, body, and environment to the conclusion that there is an extended cognitive system. From there, it moves on the conclusion that cognitive states and processes extend beyond the head into the world. In trying

to expose the fallacious nature of this two-step argument, Adams and Aizawa provisionally grant the first inference, and then challenge the second inference with the observation that “it does not follow from the fact that one has an  $X$  system that every component of the system does  $X$ ” (2009, p. 84). For instance, the components of a typical air-conditioning system include a thermostat, a compressor, an evaporation coil, and a fan. Within this system, the process of the air cooling as it passes over the evaporation coils is of the air-conditioning type. However, it would be misleading to subsume the liquefaction of Freon or the electric processes unfolding within the circuit of the thermostat under the same description as an “air-conditioning” process. Does this observation undercut our case for group cognition?

In response, we think that the analogy with machines is helpful to clarify how we think about group cognition, namely in terms of an underlying conception of *mechanistic explanation* (Craver & Bechtel, 2006). First, we identify a complex system as a machine of type  $S$  by the task that it ought to perform. This task analysis highlights a range of phenomena that have to be accomplished by the mechanism  $M$  which underlies the operation of  $S$ . Second, we identify the relevant components of  $S$  in terms of their causal relevance for the operation of  $M$ . For modular systems like an air-conditioner, this means that we break down the functions of  $M$  into a set of more basic functions, while at the same time we structurally break down  $S$  into component parts, in order to figure out what each of them does. In doing so, we then associate the more basic functions with the activities of particular components. Finally, we describe the mode of organization between the components that is required for  $M$  to work. It is thus a natural consequence of mechanistic explanation in general that the various components of a complex system of  $X$  typically do not perform functions of type  $X$ . But we do not see why this consequence would undermine an analysis of groups as cognitive systems.

We believe that the cognitive capacities of groups are amenable to the same type of mechanistic explanation as the cognitive capacity of individuals (or their brains), except for an important proviso. For most explanatory purposes, groups can be structurally decomposed into individuals and their modes of social organization. This means that unlike for individual cognition, the mechanisms of group cognition involve the activities of interacting components which have a higher degree of cognitive complexity than the system of which they are parts. For instance, information-theoretic analyses have shown that the neural connectivity of the mammalian cerebral cortex affords a high level of complexity, measured as the co-expression of functional segregation and integration (Tononi, Sporns, & Edelman, 1994). The neural architecture of our brains thus promotes high values of information integration (Sporns & Tononi, 2003) at the individual-level which are typically not sustained by the lower-bandwidth couplings between people in groups. The situation is arguably differ-

ent for eusocial insects (e.g., ants, bees) whose high degrees of co-dependence and functional integration warrant the term “super-organism” (Wheeler, 1920). What makes the attribution of cognitive capacities – such as perception, planning, and decision-making – to “super-organismic” groups so intuitively compelling is precisely the stark discontinuity between the complex collective behavior that we observe and the rudimentary cognitive resources of individual members (Couzin, 2009).

However, it remains equally true of human beings that an organized collection of parts can have more complex cognitive *capacities* than any of its parts has, despite the fact that the cognitive complexity of those *parts* outstrips the cognitive complexity of the systemically integrated whole. As an example of the ways in which the relatively low-bandwidth couplings between human beings can nevertheless give rise to emergent cognitive abilities, let us consider a group which collaboratively enacts a transactive memory system (TMS), e.g., one that is specialized for assembly tasks. Following the literature described earlier, we can associate the operation of a TMS with the level of memory differentiation, task credibility, and organized practices of retrieval coordination. Is there anything that follows from this attribution which would qualify as a systemic coupling-constitution fallacy?

First, no individual member of the group is itself a TMS, or performs in solo any of the functions that are performed by the TMS in toto. Nevertheless, every individual performs activities which causally contribute to the operation of the TMS. Some of these activities are individual memory states or processes. Other activities, such as a better acceptance of procedural suggestions, or the display of interactive cuing strategies, are not equally part of individual memory processes. However, social interactions and communication processes *between* people assume particular cognitive functions within the organization of a group that are constitutive of the operation of a TMS. For instance, transactive memory procedures are examples of group-cognitive processes that are realized only by the interactions of two or more people. Second, it is also not true that every member of the group performs *only* activities which support the TMS, or that we haphazardly lump together any kind of group interaction under the description “transactive memory procedure.” For instance, the above-cited study by Liang et al. (1995) also examined the influence of group cohesion (i.e. the level of inter-personal attraction) and social identity (i.e. the tendency of individuals to perceive themselves as members of the group). It was found that group cohesion was not positively correlated with work performance at all, and social identity, while correlated with better performance, did not mediate the impact of group training on work performance, unlike TMS did. Consequently, the individual activities and communicative interactions which were taken to constitute group cohesion and social identity are not TMS-procedures, although they were performed by components of the TMS.

Either way, our claim that groups can instantiate a TMS is not impugned by Adams and Aizawa’s criticism, because it does not even have the particular implications which they presuppose. This further means that despite their attempt to brand it as a “systemic” coupling-constitution fallacy, it remains true that we can soundly infer from the fact that a group instantiates a TMS that the transactive memory procedures by means of which the group exercises its memory capacities are realized by the communicative processes between people, rather than inside anybody’s head.

## 5. Conclusion

In this paper, we have approached the phenomenon of group cognition from the perspective of the “extended mind” thesis, as a special case of the more general claim that cognition is distributed between brain, body, and the enabling web of socio-cultural and technological scaffoldings that people inhabit. It seems to us that recent developments of the idea that minds can extend beyond a person’s skull have predominantly focused on the interactions between solitary individuals, their material environment, and the tools that they use. This has led to a relative neglect of the cognitive significance of interactions among people. However, many highly prized activities in our species are accomplished only when people think and act in groups. Hence it is tempting to consider the question whether groups can think, and what they are capable of thinking. Given the importance of people to people, it is certainly no coincidence that historically speaking, this question piqued philosophical curiosity long before our growing dependence on increasingly smart tools did the same. Following the lead of writers such as Rousseau (1762) and Hegel (1807), the “group mind” thesis became a popular idea in the conceptual landscape of the late 19th and early 20th centuries (Wegner et al., 1985). The powerful connotations of agency conjured by mentalistic concepts resonated well in an intellectual climate where groups became prominent subjects of scientific investigation, in particular in the French “collective psychology” tradition which played a foundational role in the social sciences (Wilson, 2004, Chapter 11; Wilson, 2005). However, the location and causal efficacy of the “group mind” proved to be an elusive target, because it remained unclear how the “group mind” was supposed to interact with the individuals who comprise the group (Theiner & O’Connor, 2010).

When we claim that groups can constitute cognitive systems in their own right, as we have argued, we have no desire to revive the grandeur of the traditional “group mind” thesis. In fact, a major goal of our paper has been to offer a procedure for transforming the notoriously abstract question of whether groups can think into an empirical hypothesis that admits of relatively straightforward answers. This procedure consists, roughly speaking, of three steps. First, instead of pondering about “the mark of the mental” – a unnecessarily problem-creating project that is only insignificantly altered by calling it “mark of



the cognitive” – 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-theoretical understanding of the distinctive roles which they play in the production of intelligent forms of behavior, as long as we do not mistake this classification for an explanation. Second, we employ a cluster of successful explanatory strategies in cognitive science to explain what these capacities are good for, and how they work – including a design specification of the functions they have to perform, the mechanisms by which these functions can be achieved, the principles by which those mechanisms operate. Third, we determine the extent to which these capacities are manifested by individuals and groups, how well, and how their concrete implementation differs in each case (see also Hutchins, 1995). We think that anybody who chooses not to abide by this procedure – i.e. anybody who is a “demergentist” about cognition – owes us a reason for doing so. In particular, it is incumbent on the demergentist to say what special ingredient of “the cognitive” makes the level of the individual organism metaphysically privileged to host cognitive properties, and explain why that is the case.

Keeping with our general theme of the “extended mind” thesis, we would like to end our discussion on a speculative note about the aforementioned principle of parity that was introduced by Clark and Chalmers (1998). While there has been considerable debate over the role of parity arguments for extended cognition, comparatively little attention has been paid to the individualist bias that its original formulation implies (but see Gallagher & Crisafi (2009) and Theiner (2008) for discussion). The parity principle trades on the assumption that cognitive processes in the brain of the individual provide the gold-standard which sets the bar for any other processes that also ought to count as such. However, there are many forms of advanced cognitive activities which are practically never, perhaps not even possibly, carried out inside the head of a single individual. A familiar example is our engagement in sophisticated notation-driven forms of numerical reasoning that we always carry out on paper (Dehaene, 1997). More importantly, many paradigmatic cognitive operations far too complex to be entrusted to single individuals – even when they are equipped with artifacts – are invariably carried out by teams, organizations, or entire institutions. Examples include the complex decision-making processes of large corporations, the institutional memory of the legal system, and large-scale networks of scientific collaboration. In many of these cases, we acknowledge the cognitive status of the relevant capacities by describing the task which they are meant to perform in terms that are familiar to us from the realm of individual cognition (e.g., “memory recall” or “judgment formation”). However, how do we know that there are not other, epistemically basic capacities which give rise to patterns of collective intelligence that resist an easy assimilation into the basic categories recognized by

our individual-centered taxonomies? Would we be equally inclined to consider them as cognitive capacities? Since we are not demergentists, we are not metaphysically troubled by the possibility of cognitive capacities that are, as a contingent matter, realized only at the group level. But they raise an epistemological worry: if we have no implicit comparison with individuals at all, how would we go about identifying them in the first place?

### Acknowledgments

The authors wish to thank Todd Gureckis, Leslie Marsh, Michael Roberts, Kari Theurer, Patrick Rogers, Rob Wilson, and an anonymous reviewer for *Cognitive Systems Research* for helpful suggestions on this work. This work was funded by a Killam postdoctoral fellowship from the University of Alberta granted to the first author and National Science Foundation REESE Grant 0910218 to the third author. More information about our research groups can be found at <http://ualberta.academia.edu/GeorgTheiner>, <http://cognitrn.psych.indiana.edu> and <http://mypage.iu.edu/~colallen/>.

### References

- Adams, F., & Aizawa, K. (2008). *The bounds of cognition*. Oxford: Blackwell.
- Adams, F., & Aizawa, K. (2009). Why the mind is still in the head. In P. Robbins & M. Aydede (Eds.), *The Cambridge handbook of situated cognition* (pp. 78–95). Cambridge: Cambridge University Press.
- Adams, F., & Aizawa, K. (2010). Defending the bounds of cognition. In R. Menary (Ed.), *The extended mind* (pp. 67–80). Cambridge, MA: MIT Press.
- Allen, C. (2002). Real traits, real functions? In A. Ariew, R. Cummins, & M. Perlman (Eds.), *Functions: New essays in the philosophy of psychology and biology* (pp. 373–389). New York: Oxford University Press.
- Andersson, J., & Roennberg, J. (1996). Collaboration and memory: Effects of dyadic retrieval on different memory tasks. *Applied Cognitive Psychology*, 10, 171–181.
- Argote, L. (1999). *Organizational learning: Creating, retaining and transferring knowledge*. Norwell: Kluwer.
- Ball, P. (1999). *The self-made tapestry*. Oxford: Oxford University Press.
- Baden, B. H., Baden, D. R., & Henry, S. (2000). Costs and benefits of collaborative remembering. *Applied Cognitive Psychology*, 14, 497–507.
- Bechtel, W., & Mundale, J. (1999). Multiple realizability revisited: Linking cognitive and neural states. *Philosophy of Science*, 66, 175–207.
- Block, N. (1997). Anti-reductionism slaps back. In J. Tomberlin (Ed.), *Philosophical perspectives 11: Mind, causation, and world*. Boston: Blackwell.
- Cacioppo, J. T., & Patrick, W. (2008). *Loneliness: Human nature and the need for social connection*. New York: W.W.H. Norton & Co..
- Clark, A. (2003). *Natural-born cyborgs: Minds, technologies, and the future of human intelligence*. New York: Oxford University Press.
- Clark, A. (2005). Intrinsic content, active memory and the extended mind. *Analysis*, 65, 1–11.
- Clark, A. (2008). *Supersizing the mind: Embodiment, action, and cognitive extension*. New York: Oxford University Press.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19.
- Collins, B. E., & Guetzkow, H. (1964). *A social psychology of group processes for decision-making*. New York: Wiley.
- Couzin, I. D. (2009). Collective cognition in animal groups. *Trends in Cognitive Sciences*, 13(1), 36–43.

- Couzin, I. D., Krause, J., Franks, N. R., & Levin, S. A. (2007). Effective leadership and decision making in animal groups on the move. *Nature*, 433, 513–516.
- Craver, C. F., & Bechtel, W. (2006). Mechanism. In S. Sarkar & J. Pfeifer (Eds.), *Philosophy of science: An encyclopedia* (pp. 469–478). New York: Routledge.
- Dawkins, R. (1976). *The selfish gene*. Oxford: Oxford University Press.
- Dehaene, S. (1997). *The number sense: How the mind creates mathematics*. New York: Oxford University Press.
- Diehl, M., & Stroebe, W. (1991). Productivity loss in brainstorming groups: Tracking down the blocking effect. *Journal of Personality & Social Psychology*, 61, 392–403.
- Domes, G., Heinrichs, M., Michel, A., Berger, C., & Herpertz, S. C. (2007). Oxytocin improves “mind-reading” in humans. *Biological Psychiatry*, 61, 731–733.
- Donald, M. (1991). *Origins of the modern mind: Three stages in the evolution of culture and cognition*. Cambridge: Harvard University Press.
- Dorigo, M., Bonabeau, E., & Theraulaz, G. (2000). Ant algorithms and stigmergy. *Future Generation Computer Systems*, 16, 851–871.
- Edwards, S. C., & Pratt, S. C. (2009). Rationality in collective decision-making by ant colonies. *Proceedings of the Royal Society of London, B*, 276, 3655–3661.
- Fodor, J. (1974). Special sciences (or: The Disunity of Science as a Working Hypothesis). *Synthese: An International Journal for Epistemology, Methodology and Philosophy of Science*, 28, 97–115.
- Gallagher, S., & Crisafi, A. (2009). Mental institutions. *Topoi*, 28, 45–51.
- Gilbert, M. (1989). *On social facts*. London: Routledge.
- Goldstone, R. L., & Gureckis, T. M. (2009). Collective behavior. *Topics in Cognitive Science*, 1, 412–438.
- Goldstone, R. L., Jones, A., & Roberts, M. E. (2006). Group path formation. *IEEE Transactions on System, Man, and Cybernetics, Part A*, 36, 611–620.
- Goldstone, R. L., & Roberts, M. E. (2006). Self-organized trail systems in groups of humans. *Complexity*, 11, 43–50.
- Gureckis, T. M., & Goldstone, R. L. (2006). Thinking in groups. *Pragmatics and Cognition*, 14, 293–311.
- Haidt, J., Seder, P., & Kesebir, S. (2008). Hive psychology, happiness, and public policy. *Journal of Legal Studies*, 37, S133–156.
- Hardin, G. (1968). The tragedy of the commons. *Science*, 162, 1243–1248.
- Harris, C. B., Keil, P. G., Sutton, J., & Barnier, A. J. (2010). Collaborative remembering: When can remembering with others be beneficial? In W. Christensen, E. Schier, & J. Sutton (Eds.), *ASCS09: Proceedings of the 9th conference of the Australasian society for cognitive science* (pp. 131–134). Sydney: Macquarie Centre for Cognitive Science.
- Hegel, G. W. F. (1807). *Phenomenology of spirit (A. V. Miller, Trans.)*. Oxford: Clarendon Press.
- Helbing, D., Keltsch, J., & Molnár, P. (1997). Modelling the evolution of human trail systems. *Nature*, 388, 47–50.
- Helbing, D., Schweitzer, F., Keltsch, J., & Molnár, P. (1997). Active walker model for the formation of human and animal trail systems. *Physical Review E*, 56, 2527–2539.
- Hollingshead, A. B. (1998a). Communication, learning, and retrieval in transactive memory systems. *Journal of Experimental Social Psychology*, 34, 432–442.
- Hollingshead, A. B. (1998b). Retrieval processes in transactive memory systems. *Journal of Personality and Social Psychology*, 74(3), 659–671.
- Hollingshead, A. B., & Brandon, D. P. (2003). Potential benefits of communication in transactive memory systems. *Human Communication Research*, 29(4), 607–615.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge: MIT Press.
- Johnansson, O., Andersson, J., & Roennberg, J. (2000). Do elderly couples have a better prospective memory than elderly people when they collaborate? *Applied Cognitive Psychology*, 14, 121–133.
- Kearns, M., Suri, S., & Montfort, N. (2006). An experimental study of the coloring problem on human subject networks. *Science*, 313, 824–827.
- Kosfeld, M., Heinrichs, M., Zak, P. J., Fischbacher, U., & Fehr, E. (2005). Oxytocin increases trust in humans. *Nature*, 435, 673–676.
- Le Bon, G. (1895/1960). *The crowd. A study of the popular mind*. New York: Viking Press.
- Lerner, J., Tirole, J., & Pathak, P. (2006). The dynamics of open source contributors. *American Economic Review Papers and Proceedings*, 96, 114–118.
- Lewis, K. (2003). Measuring transactive memory systems in the field: Scale development and validation. *Journal of Applied Psychology*, 88(4), 587–604.
- Lewis, K., Lange, D., & Gillis, L. (2005). Transactive memory systems, learning, and learning transfer. *Organization Science*, 16(6), 581–598.
- Liang, D. W., Moreland, R. L., & Argote, L. (1995). Group versus individual training and group performance: The mediating role of transactive memory. *Personality and Social Psychology Bulletin*, 21, 384–393.
- Lorge, I., & Solomon, H. (1955). Two models of group behavior in the solution of eureka-type problems. *Psychometrika*, 20, 139–148.
- Margulis, L. (1970). *Origin of eukaryotic cells*. New Haven: Yale University Press.
- Marsh, L., & Onof, C. (2008). Stigmergic epistemology, stigmergic cognition. *Cognitive Systems Research*, 9, 136–149.
- Mason, W. A., Jones, A., & Goldstone, R. L. (2008). Propagation of innovations in networked groups. *Journal of Experimental Psychology: General*, 137, 422–433.
- Moreland, R. L. (1999). Transactive memory: Learning who knows what in work groups and organizations. In L. Thompson, J. Levine, & D. Messick (Eds.), *Shared cognition in organizations: The management of knowledge*. Mahwah: Erlbaum.
- Moreland, R. L., Swanenburg, K. L., Flagg, J. J., & Fetterman, J. D. (2010). Transactive memory and technology in work groups and organizations. In B. Ertl (Ed.), *E-collaborative knowledge construction: Learning from computer-supported and virtual environments* (pp. 244–274). Geneva: IGI Global Press.
- Moussaid, M., Garnier, S., Theraulaz, G., & Helbing, D. (2009). Collective information processing and pattern formation in swarms, flocks, and crowds. *Topics in Cognitive Science*, 1, 469–497.
- Munn, C. A. (1986). The deceptive use of alarm calls by sentinel species in mixed-species flocks of neotropical birds. In R. W. Mitchell & N. S. Thompson (Eds.), *Deception: Perspectives on human and nonhuman deceit* (pp. 169–176). Albany: SUNY Press.
- Norman, D. (1991). Cognitive artifacts. In J. M. Carroll (Ed.), *Designing interaction: Psychology at the human-computer interface*. Cambridge: Cambridge University Press (p. ix, 333 p.).
- Page, S. E. (2007). *The difference: How the power of diversity creates better groups, firms, schools and societies*. Princeton: Princeton University Press.
- Pavitt, C. (2003). Colloquy: Do interacting groups perform better than aggregates of individuals? Why we have to be reductionists about group memory. *Human Communication Research*, 29(4), 592–599.
- Propp, K. (1999). Collective information processing in groups. In L. Frey, D. S. Gouran, & M. S. Poole (Eds.), *Handbook of group communication theory and research* (pp. 225–250). Thousand Oaks: Sage.
- Propp, K. (2003). In search of the assembly bonus effect: Continued exploration of communication's role in group memory. *Human Communication Research*, 29(4), 600–606.
- Putnam, R. D. (2001). *Bowling alone: The collapse and revival of American community*. New York: Simon & Schuster.
- Rawls, J. (2001). *Justice as fairness: A restatement*. Cambridge, MA: Belknap Press.
- Rehmeyer, J. J. (2007). Mapping a medusa: The Internet spreads its tentacles. *Science News*, 171(23), 387–388.
- Resnick, M. (1994). *Turtles, termites, and traffic jams: Explorations in massively parallel microworlds*. Cambridge, MA: MIT Press.
- Roberts, M. E., & Goldstone, R. L. (2009). *Adaptive group coordination proceedings of the thirty-first annual conference of the cognitive science society* (pp. 2698–2704). Amsterdam, Netherlands: Cognitive Science Society.
- Roberts, M. E., & Goldstone, R. L. (submitted for publication). Adaptive group coordination and role differentiation. *International Journal of the Commons*.

- Rousseau, J.-J. (1762). *The social contract* (M. Cranston, Trans.). London: Penguin Books.
- Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *The Journal of Philosophy*, 101, 389–428.
- Simon, H. (1969). *The sciences of the artificial*. Cambridge, MA: MIT Press.
- Smith, A. (1776). *An inquiry into the nature and causes of the wealth of nations*. London: Methuen and Co..
- Smith, E. R., Seger, C. R., & Mackie, D. M. (2007). Can emotions be truly group-level? Evidence regarding four conceptual criteria. *Journal of Personality and Social Psychology*, 93, 431–446.
- Sporns, O., & Tononi, G. (2003). Measuring information integration. *BMC Neuroscience*, 4(31).
- Steiner, I. D. (1972). *Group process and productivity*. New York: Academic Press.
- Theiner, G. (2008). From extended minds to group minds: Rethinking the boundaries of the mental. *Doctoral dissertation, Department of Philosophy & Cognitive Science Program, Indiana University, Bloomington*.
- Theiner, G. (2010). Making sense of group cognition: The curious case of transactive memory systems. In W. Christensen, E. Schier, & J. Sutton (Eds.), *ASCS09: Proceedings of the 9th conference of the Australasian society for cognitive science* (pp. 334–342). Sydney: Macquarie Centre for Cognitive Science.
- Theiner, G., & O'Connor, T. (2010). The emergence of group cognition. In A. Corradini & T. O'Connor (Eds.), *Emergence in science and philosophy* (pp. 78–117). New York: Routledge.
- Tollefsen, D. (2004). Collective epistemic agency. *Southwest Philosophy Review*, 20(1), 55–66.
- Tollefsen, D. (2006). From extended mind to collective mind. *Cognitive Systems Research*, 7(2–3), 140–150.
- Tononi, G., Sporns, O., & Edelman, G. M. (1994). A measure for brain complexity: Relating functional segregation and integration in the nervous system. *Proceedings of the National Academy of Sciences Neurobiology*, 91(11), 5033–5037.
- Walsh, J. P., & Ungson, G. R. (1991). Organizational memory. *Academy of Management Review*, 16, 57–91.
- Wegner, D. M. (1986). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of group behavior* (pp. 185–208). New York: Springer Verlag.
- Wegner, D. M. (1995). A computer network model of human transactive memory. *Social Cognition*, 13, 319–339.
- Wegner, D. M., Erber, R., & Raymond, P. (1991). Transactive memory in close relationships. *Journal of Personality and Social Psychology*, 61, 923–929.
- Wegner, D. M., Giuliano, T., & Hertel, P. (1985). Cognitive interdependence in close relationships. In W. J. Ickes (Ed.), *Compatible and incompatible relationships* (pp. 253–276). New York: Springer Verlag.
- Weldon, M. S., & Bellinger, K. D. (1997). Collective memory: Collaborative and individual processes in remembering. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(5), 1160–1175.
- Wheeler, G. (1920). The termitodoxa, or biology and society. *Scientific Monthly*, 10, 113–124.
- Wilensky, U. (1997). NetLogo traffic basic model. Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University. <<http://ccl.northwestern.edu/netlogo/models/TrafficBasic>>.
- Wilson, R. (2004). *Boundaries of the mind: The individual in the fragile sciences*. Cambridge: Cambridge University Press.
- Wilson, R. (2005). Collective memory, group minds, and the extended mind thesis. *Cognitive Processes*, 6(4), 227–236.
- Wilson, R. (in press). Extended vision. In N. Gangopadhyay, M. Madary & F. Spicer (Eds.), *Perception, action and consciousness*. Oxford: Oxford University Press.
- Wimsatt, W. C. (1986). Forms of aggregativity. In M. G. Grene, A. Donagan, A. N. Perovich, & M. V. Wedin (Eds.), *Human nature and natural knowledge* (pp. 259–291). Dordrecht: Reidel.
- Wittenbaum, G. M. (2003). Putting communication into the study of group memory. *Human Communication Research*, 29(4), 616–623.
- Wright, R. (2001). *Nonzero*. New York: Vintage.