Extended Minds, Collective Minds, and Group Agents

By Tertia Davis

Submitted to
Central European University
Department of Philosophy

In partial fulfilment of the requirements for the degree of Master of Philosophy

Supervisor: Professor Simon Rippon

Budapest, Hungary
2020

Abstract

We often speak of groups as if they are agents or minds in their own right. I do not think this is purely metaphorical, and that sometimes we pick out a distinct group agent that is capable of intentions, actions, thoughts, beliefs, desires, and other mental states that do not reduce to the aggregate mental states of the individual human beings that compose the group. In this thesis I make an argument for a robust, realist conception of the group agent based upon the Hypothesis of Extended Mind (HEM).[1] Using a central principle of this theory, I claim that the mind extends, not only into the physical environment, but socially as well. This socially extended mind provides the basis of group mental states, which function similarly to those of individual human minds, specifically as distributed cognitive systems. I then address the primary objection to this argument, which is that the robust group agent entails an impossible group mind. But instead of rejecting the conclusion that a strong conception of the group agent implies a potentially conscious group or collective mind, I argue that this is not necessarily an unreasonable or undesirable outcome.

[1] Andy Clark and David Chalmers, "The Extended Mind," *Analysis* 58, no. 1 (1998): 7–19; Andy Clark, *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*, Philosophy of Mind (Oxford: Oxford University Press, 2008).

Acknowledgements

I would like to express my gratitude to the CEU Department of Philosophy for providing me with the opportunity to study with so many knowledgeable, thoughtful people. My professors and cohort have consistently challenged, supported, and inspired me. In particular, I greatly appreciate all of the help and feedback from my thesis supervisor Simon Rippon. Thank you to Loren who encouraged me from the very beginning, and to my parents, Danny, and Eislin, who made this all possible. I am deeply thankful for my time at CEU and the many kind friends I made during my time in Budapest.

Dedicated to my friend Raza Khan.

Table of Contents

Introduction

> …it can make sense to think there exists, inside your brain, a society of different minds. Like members of a family, the different minds can work together to help each other, each still having its own mental experiences that the others never know about. Several such agencies could have many agents in common, yet still have no more sense of each other's interior activities than do people whose apartments share opposite sides of the same walls. Like tenants in a rooming house, the processes that share your brain need not share one another's mental lives.[2]

The central debate in the field of social ontology concerns the reducibility of collective intentions to the intentions of the individuals that compose the group, and the related question of the ontological status of the group agent. In this thesis, I make an argument for a robust, realist conception of the group agent based upon the Hypothesis of Extended Mind (HEM).[3] As a realist about group agents, I think that when we say something like "the court ruled in favor of the defendant," or "Nestle is responsible for environmental devastation," we may pick out a distinct group agent that is capable of intentions, actions, thoughts, beliefs, desires, and moral responsibility in a way that is irreducible to an aggregate of individuals. This helps us to make sense of the fact that groups frequently act in ways that do not reflect the intentions of the members, and even behave as organisms with their own autonomous, integrated, and rational point of view, including goals for their own persistence, that are fully distinct from those of their constituent parts.

Some accounts of collective intentionality assume that our reference to groups as agents is only metaphorical, or suggest that collective intentions are relatively limited in scope, operating only as a mechanism to coordinate plans between individuals. This somewhat conservative approach typically excludes many instances from meeting the definition of a collective action, like the behavior of mobs or accidental group actions. But

[2] Marvin Minsky, *The Society of Mind* (New York: Simon & Schuster, 1988), 290.
[3] Clark and Chalmers, "The Extended Mind"; Clark, *Supersizing the Mind*.

these reductive approaches to group agency, I think, are strongly founded on the desire to avoid mysterious metaphysical claims about superorganisms and group minds.

The HEM provides us with a way to approach the idea of a mind that is socially extended, or even properly collective, from a position with some empirical support. Popular physicalist, functionalist views of cognition are hard pressed not to accept Andy Clark and David Chalmers's Parity Principle,[4] which can be applied to other mental states beyond cognition, and may help us understand how consciousness could be attributed to a group. I think that theories of the extended mind and distributed cognition support the notion of a robust group agent that possesses "a mind of its own."[5] My goal is to use these theories of mind to provide further support for the realist accounts of group agents suggested by Christian List, Philip Pettit, and Hand Bernhard Schmid.[6]

In this thesis I use the central principle of the HEM to strengthen the position of realist conceptions of the group agent, which I think are mutually supportive theories, but they are not often discussed together in this manner. Overall, my project is to encourage us to take more seriously the possibility of agents and minds that do not closely resemble individual human beings, because I think the ability to identify other kinds of minds and agents could be an incredibly powerful explanatory tool in understanding the behavior of currently-existing groups, as well as providing us with strategies for discovering, recognizing, or creating, different minds or intelligences.

In the first chapter, I cover the central concepts in this field, and claim that the Parity Principle applies to the extension of the mind into the social world and provides a basis for

[4] Clark and Chalmers, "The Extended Mind."

[5] Philip Pettit, "Groups With Minds of Their Own," in *Socializing Metaphysics: The Nature of Social Reality*, ed. Frederick Schmitt (New York: Rowman & Littlefield, 2003), 167–93.

[6] Christian List and Philip Pettit, *Group Agency: The Possibility, Design, and Status of Corporate Agents* (Oxford: Oxford University Press, 2013); Hans Bernhard Schmid, "Plural Self-Awareness," *Phenomenology and the Cognitive Sciences* 13, no. 1 (March 2014): 7–24.

group mental states. In the second chapter, I claim that group mental states are distributed cognitive systems very similar to our own, and that this implies the existence of an irreducible, thinking, intentional group agent. The final chapter is dedicated to what I see as the primary objection to this argument, which is that the robust group agent entails an impossible group mind. Although I agree that a strong conception of the group agent (and the HEM) do imply a potentially conscious group or collective mind, I argue that this is not necessarily an unreasonable or undesirable result.

Chapter One: Extended Minds & Group Mental States

Collective Intentionality and Group Agency

Although I cannot hope to give a full account of the theories about collective intentionality and group agency here, I will briefly describe the different positions in order to explain why I support a robust realist conception of the group agent. The general approaches to group agency can be categorized as eliminativist/reductionist, plural subject, and realist positions. But first, what is the relationship between collective intentionality and group agency?

Generally, we use the term intentionality to speak of the mind's directedness upon objects, states of affairs, properties, or events. Intentional states are things like beliefs, desires, and thoughts that are "at" or "about" something. The intention to act, or to do something deliberately, is often what people refer to when using the word "intend," but this concept can be located within a broader theory of intentionality.[7] An intention to act is just one kind of intentionality, which is a more inclusive notion of a mind's directedness toward something (like the world, or itself). Intentionality is typically regarded as the sign of mentality, the characteristic displayed exclusively by mental states.[8] There are some significant disagreements regarding the relationship between intentionality and mental states, but most people who have adopted this terminology agree that there is a close relationship between them. There is also an important relationship between intentionality and agency, because intentions are necessary preconditions for action, at least according to the standard conception. An agent must be capable of intentionality (directedness) in order to from a specific intention to act.

[7] John R. Searle, "The Intentionality of Intention and Action," *Cognitive Science* 4, no. 1 (1980): 47–70.

[8] Tim Crane, "Intentionality as the Mark of the Mental," in *Contemporary Issues in the Philosophy of Mind*, ed. Anthony O'Hear, 1998.

In everyday language, we often say things like "The United States is preparing to invade Iraq," and "Amazon mistreats its employees." Nations, corporations, sports teams, cultural groups, families, and even couples are often spoken of as if they are a single unit, capable of feeling and acting as cohesive agent. This attribution of agency often sounds very straightforward, and we do seem to think that a country can mourn, or that a company can have moral responsibility. But we might dismiss this sort of talk as metaphorical shorthand, and claim that this manner of speaking just refers to the aggregation of the individual intentions and actions.

Eliminativists about group agents think that collective intentions can be reduced to individual intentions, and deny that there is any causal-explanatory need for the idea of a group agent or group subject.[9] Michael Bratman claims that individual intentions are distinct, irreducible attitudes that "are central to our shared understanding of ourselves as intelligent agents,"[10] and although group intentions may be understood as joint or shared in a particular way, they are just a complex state of affairs that consist only of the relations of attitudes between the individuals of a group, and are not genuinely group-level intentions. He acknowledges the significance of our sociality, but believes that all talk of group agents is, or should be, purely metaphorical, because the relevant attitudes that lead to collective intentions exist only in the minds of the individuals. Other eliminativists also deny that collective intentions are anything "over and above" the aggregation of individual intentions. Chant acknowledges the reality and importance of collective actions, but says "All the agency required for the collective intention is possessed by the individuals in the group."[11]

[9] Michael Bratman, *Shared Agency: A Planning Theory of Acting Together* (New York: Oxford University Press, 2014), 124–27.

[10] Michael Bratman, "Shared Intention," *Ethics* 104, no. 1 (1993): 97–98.

[11] Sara Rachel Chant, "Collective Action and Agency," in *The Routledge Handbook of Collective Intentionality*, ed. Marija Jankovic and Kirk Ludwig (New York: Routledge, 2018), 23.

Perhaps all of our talk of group intentions, agents, and mental states is not intended literally, and they are not all that similar to the intentions, actions, and mental states of individuals after all; maybe complex group actions can be explained by an appropriate and nuanced understating of individual intentions and interactive knowledge without over-populating our ontology. For this reason, eliminativists about group agents posit that collective intentions and actions are reducible to the constitutive members, and there is no explanatory need for the group agent.

I think the most powerful objection to reductive and aggregative accounts of group intentions and actions is provided by Christian List and Philip Pettit who analyze the discursive (or doctrinal) paradox that arises when groups attempt to aggregate rational individual judgements only to result in irrational and contradictory judgements at the group level.[12] It turns out that it isn't just majoritarian voting strategies that lead to this impossibility result, but in fact "no aggregation function will generate rational group attitudes in accordance with some minimal conditions."[13] And yet, groups clearly do behave in rational ways, making decisions and acting on beliefs, desires, and toward their continued ability to exist, as a group. If we can't make sense of the rational actions of groups in terms of the individuals that compose them, then that suggests that reductive accounts cannot explain how groups can and do act in unified and autonomous ways, leaving room for a more robust conception of collective agency and intentionality.

If we, like Pettit, take a broadly functionalist view of intentionality, understanding the intentional system as a system with beliefs and desires, that is "exposed perceptually to a

[12] Christian List and Philip Pettit, "Aggregating Sets of Judgments: An Impossibility Result," *Economics and Philosophy* 18, no. 1 (April 2002): 89–110.

[13] List and Pettit, *Group Agency*, 58.

certain sort of environment in a way that makes sense, and makes sense non-accidentally, in belief-desire terms,"[14] then:

> Intentional subjects are a species of agent. Roughly, they are agents that engage with their environment in such a way that we ascribe beliefs and desires to them; by contrast with stimulus-response automatons, they act on the basis of how they construe their situation and how they feel about it.[15]

This inclusive definition of the intentional agent potentially includes at least some non-human living creatures, and other sufficiently complex systems with the right sort of causal relevance, which I think is important so that we can better understand and look for agents unlike ourselves. The human being is the paradigmatic example of a thinking intentional subject, but Pettit allows that other systems may qualify. In fact, he proposes that not only are groups intentional agents, but that they are fully thinking agents with "minds of their own."[16] Pettit is a realist about group agents as "institutional persons," but does admit that they differ significantly from individual persons as they are much more limited in scope, spontaneity, and are "not centers of perception or memory or sentience."[17] If we take for granted that a group mind lacks those features, we may question how his notion of a group mind qualifies as a mind at all. However, I do consider this view to be more appealing than the aforementioned reductive accounts in that it addresses how groups act in ways that cannot be anticipated through the examination of individual intentions, regardless of individual plans and interactive knowledge.

Others who object to reductionism, like John Searle, often consider intentionality to be biologically primitive, because it "functions only against a background of

[14] Philip Pettit, *The Common Mind: An Essay on Psychology, Society, and Politics* (New York: Oxford University Press, 1996), 10.
[15] Pettit, 10.
[16] Pettit, "Groups With Minds of Their Own."
[17] Pettit, 188.

nonrepresentational mental capacities."[18] This goes for collective intentions as well as for individual intention, because "We-intentions cannot be analyzed into sets of I-intentions, even I-intentions supplemented with beliefs, including mutual beliefs, about the intentions of other members of a group."[19] And this suggests that groups, as well as individuals, have mental capacities. Although reductive accounts may have some advantages, particularly the ability to discuss the actions of mobs and disorganized groups in a way that may be challenging for the more robust accounts, I think that there are convincing arguments that aggregative accounts simply can't provide satisfying explanations for group intentions and actions by appeals to individual intentions and actions alone.

Another non-reductive account of shared intentions is the "plural subject" account where commitments are shared among people in such a way that they form plural subjects of experience.[20] Margaret Gilbert doesn't think that shared intentions or actions can be reduced to, or explained by, appeals to individual intentions or actions, claiming that "Among other things, a joint commitment links people together or *unifies* them in such a way that it makes sense to say that they have *created something new*, something that is more, at least, than a mere aggregate of persons."[21] This is an appealing approach to group agents, because it provides an explanatorily useful concept of the plural subject to make sense of our everyday attribution of agency to groups, capable of explaining how we plan together and function as a coherent unit. However, the more realist concepts of group agency may be challenged on the

[18] John Searle, *Intentionality: An Essay in the Philosophy of Mind* (Cambridge and New York: Cambridge University Press, 1983), viii.

[19] John Searle, *Consciousness and Language* (Cambridge: Cambridge University Press, 2002), 93.Within this paper I will be using the terms "we-intentions" and "collective intentions" and "shared intentions" interchangeably.

[20] Margaret Gilbert, *On Social Facts* (Princeton, NJ: Princeton University Press, 1992); Margaret Gilbert, *Joint Commitment: How We Make the Social World* (New York: Oxford University Press, 2014).

[21] Italics in the original, Gilbert, *Joint Commitment*, 67.

grounds that they seem to imply metaphysically dubious entities like supra-individuals, group minds, and collective consciousness.

Plural subject accounts may also be accused of sliding back into a sort of reductionism, because they also embrace methodological individualism and "attempt to account for the structure of collective intentionality without letting any genuine collectivity enter the scene. 'We-ness' is the topic, yet at the same time it is stressed that it is a feature of individuals—and not of an actual 'we.'"[22] I particularly like Hans Bernhard Schmid's robust realist account because it does not assume methodological individualism, unlike the above reductive accounts, or even Gilbert's plural subject account, which still relies on the ontological supremacy of the individual human being. In order to not beg the question against robust notions of the group agent, we cannot begin with the assumption that "society consists entirely of individuals," as Searle claims, because it seems reasonable to wonder if society is also composed of other sorts of entities. And in fact, this statement of Searle's is made in order to explicitly reject the possibility of a group mind or consciousness, foreshadowing the primary objection to my argument that I will address in the final chapter.

I think that robust accounts of group agency not only provide us with an interesting and necessary level of explanation for group behavior, they also offer a framework for understanding a variety of collective intentional states, like group emotions and states with shared affective components. In addition to mental states like collective knowledge or beliefs, Schmid thinks that we can attribute emotions to groups in a simple, straightforward sense, because collective emotions have a single token, feeling episode, or phenomenal emotional experience.[23] But the idea that a group might be the subject of an emotion that does not

[22] Hans Bernhard Schmid, *Plural Action: Essays in Philosophy and Social Science* (Dordrecht: Springer, 2009), 32.

[23] Schmid, 83; Hans Bernhard Schmid, "The Feeling of Being a Group," in *Collective Emotions*, ed. Christian von Scheve and Mikko Salmela (New York: Oxford University Press, 2014), 9.

reduce to the emotions of the individuals that compose it starts to raise many questions about whether we can ascribe some sort of consciousness or experience to such a subject, itself composed of conscious parts. As counter intuitive as the implication of a collective consciousness may sound initially, I think it is a possibility that we should thoroughly explore. And this is a problem I would like to address, because I think that aside from this challenge, realist accounts of group agents provide the most illuminating way to understand group behavior, that reductive accounts are not capable of. The more middle-of-the road accounts I also find unsatisfying, because, although they also acknowledge the weakness of aggregative accounts, they tend to rely upon the same individualistic assumptions, and when pressed, collapse either into reductionism or realism.

I think that a great deal of the opposition to the robust group agent is the inability to imagine how such a thing could be so, and how there could be an entity composed of conscious parts that are somewhat or fully unable to access the group-level mental states. I think that the HEM provides us with a clear way to envision how a group of individuals could give rise to group mind, and a thinking group agent.

Extended Minds

The HEM, as conceived of by Clark and Chalmers, is the claim that the mind is not limited by "the boundaries of skin and skull."[24] For the functionalist, there seems to be no reason to assume that minds must be restricted to biological matter like our own brains. And if we examine our assumptions, there are no logical grounds to rule out the possibility of minds composed of other materials. How do we recognize a mind or cognitive processes if we don't assume that only brains very similar to our own are capable of operating in this way? Clark and Chalmers propose the Parity Principle (PP):

[24] Clark and Chalmers, "The Extended Mind," 7.

> If, as we confront some task, a part of the world functions as a process which, were it to go on in the head, we would have no hesitation in accepting as part of the cognitive process, then that part of the world is (for that time) part of the cognitive process.[25]

This approach opens up a new world of cognitive possibility, and encourages us to look for cognitive processes external to the brain. I would like to argue that this principle applies just as well to cognitive processes that are distributed across a group as to cognitive processes that integrate the inanimate environment into the individual mind. But first I will clarify how Clark and Chalmers argue for and defend this principle, given that it is a necessary pillar of my argument.

Clark and Chalmers emphasize that it is not simply the physical basis for cognition that extends into the environment, but that the environment becomes an active part of, and plays a causal role in, a coupled system that "can be seen as a cognitive system in its own right."[26] This would be compatible with the idea that mental states are still fully internal, but they suggest that some mental states, like beliefs, are at least somewhat constituted by the environment. According their functionalist stance, wherein a belief is defined as such by the causal role it plays, external features can play the same causal function as beliefs, so they claim that "the mind extends into the world."[27]

The primary thought experiments used to defend the PP, and thus the HEM, are those of Otto and Inga. Otto has Alzheimer's disease and uses a notebook to record information that he has difficulty remembering. And not only does this notebook serve as a repository for data, it fulfills the role of memory and belief, in the exact same way that Inga's internal memory does. One way to refute this principle would be to restrict the mind to specific biological materials like our own brain, but it seems perfectly plausible that a mind could exist in other substrates, whether created by human beings or evolved in alien life forms.

[25] Clark and Chalmers, 8; Clark, *Supersizing the Mind*, 77.
[26] Clark and Chalmers, "The Extended Mind," 8.
[27] Clark and Chalmers, 12.

However, if we accept the PP, as I think we should, this raises many questions about where to set the boundaries of the mind, if not at the border of the individual brain and body? It seems as if their argument could easily lead to a conception of mind that is infinite, or so overly extended as to be nearly meaningless. Clark and Chalmers propose three conditions that a system must meet in order to be understood as cognitively coupled: high levels of trust, reliability, and accessibility.[28] This allows us to distinguish between a desktop computer, which presumably would not meet these conditions, and a cane used by someone who is visually impaired, that serves as an integrated, trusted, reliable, and accessible extension of their mind, just as our other appendages do.

Some theorists who wish to reject the conclusion of the HEM, like Brie Gertler, do so by refuting other premises of Clark and Chalmers's arguments, because she does not find fault with the PP. She claims that if standing states or nonconscious cognitive processes are considered to be part of the mind, then the mind does extend, but to an excessive and implausible degree, potentially including external devices like notebooks and computers, and even the minds of other people.[29] This leads her to adopt a very narrow definition of mind such that it consists only of occurrent states, the experience of a stream of consciousness, and not standing states like beliefs and desires. However, the decision to exclude these standing states from our picture of the cognitive and mental seems to be a position no less extreme than postulating that the mind extends prolifically. I think that denying that standing states are part of the mind is an even less desirable conclusion than accepting a high degree of mind extension.

[28] Clark and Chalmers, 17.

[29] Brie Gertler, "Overextending the Mind?," in *Arguing About the Mind*, ed. Brie Gertler and Lawrence Shapiro (New York: Routledge, 2007), 192–206.

It has been suggested that the PP applies not only to cognitive processes and mental states, but to consciousness itself.[30] However, I would like to make a slightly more modest claim myself, and just apply the PP as it is to include the social world as well as the physical world. Clark and Chalmers reject the postulation that consciousness extends,[31] but do recognize that, according to the HEM, mental states like beliefs could "be embodied in one's secretary, one's accountant, or one's collaborator."[32] In the next subsection, I will discuss the application of the PP to the social world, and the relevance of this claim in support of the robust group agent.

Socially Extended Minds

One of the most interesting implications of the HEM is the possibility that the mind extends, not only into the inanimate environment, but socially, to include the minds of others. Clark and Chalmers attribute the coupling process in these cases to the use of language to facilitate communication between people. But we have a high of capacity to communicate beyond spoken and written language, and should consider that:

> 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.[33]

Clark and Chalmers suggest that minds extend socially in a limited way, with language as the only mediator for this coupling process. But there are many other channels for communication and mutual understanding than through spoken or written language alone. Other examples of potential social coupling mechanisms include body posture, facial

[30] Karina Vold, "The Parity Argument for Extended Consciousness," *Journal of Consciousness Studies* 22, no. 3–4 (2015): 16–33.

[31] Clark and Chalmers, "The Extended Mind," 10.

[32] Clark and Chalmers, 18.

[33] Georg Theiner, Colin Allen, and Robert L. Goldstone, "Recognizing Group Cognition," *Cognitive Systems Research*, Special Issue on Extended Mind, 11, no. 4 (December 1, 2010): 380.

expression, cooperative action, mind and behavior reading, joint attention, and social norms, among others.[34] And, unlike notebooks and other artifacts, human minds are similarly active, with high levels of processing speed and bandwidth, and varying degrees of shared epistemic background. Indeed, other minds seem a likely place to look for an extended mind.

According to the PP, if some process happens outside of the brain that would be considered cognitive if it was an internal process, then it should still be considered cognitive. The PP applies to the concept of mental states similarly, so that if the environment plays the role of a state that would be considered mental if it was internal, then it should be considered mental as well, even if it is not specifically in the brain. The environment that we speak of may consist of tools or artifacts, but also other human beings. By this formulation of the social PP, if some function is fulfilled by a suitably coupled pair or group of people that would be considered mental if it was internal, it should likewise be considered mental, even if part of the process is external. In this way, the mind can extend to encompass more than one individual as proper parts.

We might also consider that some mental states are, at times, even entirely external, like the idea that legal, educational, and cultural institutions are "mental institutions," that enable, or even fully constitute certain cognitive processes.[35] Not only do these social systems shape what we can think or imagine, but a legal agreement like a contract "is in some real sense an expression of several minds externalized and extended into the world, instantiating in external memory an agreed-upon decision, adding to a system of rights and laws that transcend the particularities of any individual's mind."[36] By this reasoning, there are some

[34] Holger Lyre, "Socially Extended Cognition and Shared Intentionality," *Frontiers in Psychology* 9 (May 28, 2018): 6.

[35] Shaun Gallagher and Anthony Crisafi, "Mental Institutions," *Topoi* 28, no. 1 (March 2009): 45–51; Shaun Gallagher, "The Socially Extended Mind," *Cognitive Systems Research* 25–26 (December 2013): 4–12.

[36] Gallagher, "The Socially Extended Mind," 6.

clearly cognitive processes (consideration, judgement, and so forth) that happen, not only within individual minds alone, but even entirely external to the individual minds that compose a group. This sort of socially extended cognition might occur in the "we-space," mediated by our various gestures and expressions in such a way that we directly experience and participate in cognitive processes with others in a nonconceptual manner.[37] We might also think that we are able to directly perceive the mental states of others on some occasions,[38] so that the connection between our mental states is unmediated and truly shared.

These views of socially extended cognition are non-reductive, positing cognitive processes as a particular kind of dynamic relationship that cannot be analyzed at the level of the components alone. Looking at socially extended cognition in this way also satisfies Gilbert's "disjunctive criterion," whereby individual intentions within the group need not correlate to the shared intention.[39] If minds are socially extended, we can see how groups can make decisions that aren't reflected by the desires or intentions of the individual members. A mind that is socially extended, a group or collective mind, can in fact behave in ways that do not correspond with, or may directly oppose the wishes, reasoning, and best interests of the group members.

As an example of the socially extended mind, Krueger discusses the cognitive development of infants. The physical interventions of caregivers in their interactions with a baby, like gesture, expression, and so forth, serve to "regulate attention and emotion," and thus "are part of the infant's socially extended mind; they are external mechanisms that

[37] Joel Krueger, "Extended Cognition and the Space of Social Interaction," *Consciousness and Cognition* 20, no. 3 (September 1, 2011): 643–57.

[38] Shaun Gallagher, "Direct Perception in the Intersubjective Context," *Consciousness and Cognition*, Social Cognition, Emotion, and Self-Consciousness, 17, no. 2 (June 1, 2008): 535–43.

[39] Margaret Gilbert, "Shared Intention and Personal Intentions," *Philosophical Studies* 144, no. 1 (2009): 172.

enable the infant to do things she could not otherwise do, cognitively."[40] The infant-caregiver relationship is particularly interesting to me. Corporations, states, and dyads of people moving furniture are common ways to illustrate collective action, mental states, and group agency, but certainly there are closer and more sensitive relationships that may have more substantial cognitive coupling. Donald Winnicott, an early twentieth-century psychoanalyst, is known for his claim that "'There is no such thing as a baby' –meaning that if you set out to describe a baby, you will find you are describing a *baby and someone.* A baby cannot exist alone, but is essentially part of a relationship."[41] The individual, it seems, is not necessarily prior to or more fundamental than "us."

Instead of focusing strictly on cases of shared agency in terms of helping to plan collective actions, we should also give attention to other kinds of relationships and manifestations of extended cognition and shared mental states. Infants are incredibly sensitive to caregiver proximity, attention, touch, facial expressions, eye contact, and language. The heart rates, breathing, and temperature of newborns and caregivers synchronize during skin-to-skin contact, and many neurophysiological processes are regulated through caregiver interaction.[42] Babies are entirely vulnerably and dependent on the resources, efforts, and attention of caregivers for a significant period of time, during which it is difficult to pinpoint exactly when we can claim they become conscious beings. The stability of levels of oxytocin, the primary hormone responsible for social bonding and reciprocity, may be consistent over

[40] Joel Krueger, "Ontogenesis of the Socially Extended Mind," *Cognitive Systems Research*, Socially Extended Cognition, 25–26 (December 1, 2013): 40.

[41] Donald W. Winnicott, *The Child, the Family, and the Outside World* (Cambridge, MA: Perseus Publishing, 1987), 88. Italics in original.

[42] Ruth Feldman et al., "Mother and Infant Coordinate Heart Rhythms through Episodes of Interaction Synchrony," *Infant Behavior and Development* 34, no. 4 (December 2011): 569–77; Myron A. Hofer, "Hidden Regulators in Attachment, Separation, and Loss," *Monographs of the Society for Research in Child Development* 59, no. 2/3 (1994): 192–207.

time and transferred cross-generationally through caregiving behaviors,[43] further indicating how the mental and emotional processes of an individual are scaffolded, or may even be considered an extension of caregiver mental processes.

Winnicott also introduced the mother-baby unit to discussions of child development, but this sentiment is also popular among various birth communities, the basic idea being that the mother (or other full-time caregiver) are so connected, physically, mentally, and emotionally, during the first year in particular, that they constitute a single unit or organism. The intense bonding and cognitive development in that period I think shows the sort of intimacy required for a "phenomenological fusion of feelings,"[44] so that the infant and sensitive caregiver may have a genuinely collective intentional state.

The methodological individualist approach to social ontology presumes the metaphysical primacy and priority of the individual human being, and that "we" is just a collection of already established individual "I's," but I think this assumption should be challenged. Perhaps "I" and "we" arise simultaneously, or the collective sense of "us" even precedes and provides the foundation for the concept of "I" and self-awareness, as Schmid suggests.[45] And don't think that this is merely armchair speculation on the part of philosophers; scientists in many fields are also challenging the idea that single organisms are the fundamental kind of individual.[46] Perhaps individuality is only something that we acquire through certain social interactions and background capacities.

[43] Ruth Feldman et al., "Parental Oxytocin and Early Caregiving Jointly Shape Children's Oxytocin Response and Social Reciprocity," *Neuropsychopharmacology* 38, no. 7 (June 2013): 1154–62.

[44] Schmid, *Plural Action: Essays in Philosophy and Social Science*, 80.

[45] Schmid, "Plural Self-Awareness," 23.

[46] Frédéric Bouchard and Philippe Huneman, eds., *From Groups to Individuals: Evolution and Emerging Individuality*, Vienna Series in Theoretical Biology (Cambridge, MA: The MIT Press, 2013).

Not only does the PP theoretically support the idea that cognition can extend into groups of people, but there is also a variety of evidence from other fields, particularly studies of infant cognitive development, that groups of people do constitute complex cognitive systems in their own right. Groups are capable of mental functions like judgement, memory, emotion, and mental states that are at least partially external and inexplicable in terms of the mental states of the constituent members. Even if you accept this, you might argue that a collective mental state is still simply too different from our own individual experiences of these mental states in their phenomenal character, and that even if they are cognitive systems, they are not agents in the same way that human beings are. Having established that the PP is applicable to socially extended cognition, resulting in a strong notion of collective mental states, the next chapter is dedicated to showing that the existence of a collective mental state indicates a strong notion of a group agent, in possession of a mind in a meaningful sense, whose mental processing and exercise of agency closely resembles our own.

Chapter Two: Group Mental States & Group Agents

Distributed Cognition and Soft Ecological Control

In the previous section, I used the PP to establish that mental states, and thus minds, can extend socially, and provide the foundation for truly collective mental states. Now I will argue that these socially extended, group mental states are distributed cognitive systems, like our own minds, and that these mental capacities allow for a group-level agent that acts upon the world as an integrated unit, just as we do.

Prior to the development of the HEM by Clark and Chalmers, anthropologist Edwin Hutchins articulated an influential theory of distributed cognition.[47] His research is founded in a strong tradition in the social sciences and psychology, and details an account of naval navigational systems. He shows that they function as a single cognitive, or computational, unit that displays characteristics and capacities not possessed by the component members. The naval ship is one example of how cognitive systems can be extended across large numbers of people and artifacts, but other possibilities for distributed cognitive systems proliferate across scientific disciplines, from the study of slime mold and cockroaches, to CERN, and the Hubble Space Telescope.[48] It is important to note that discussion of distributed cognitive systems within the social and cognitive sciences have included large group-level systems of humans and many other living organisms (that may or may not include non-living objects from the environment), as well approaching the individual mind as a similarly distributed system.

[47] Edwin Hutchins, *Cognition in the Wild* (Cambridge, MA: MIT Press, 1995).

[48] Andy Clark, *Being There: Putting Brain, Body, and World Together Again* (Cambridge, MA: MIT Press, 2001); Karin Knorr Cetina, *Epistemic Cultures: How the Sciences Make Knowledge* (Cambridge, MA: Harvard University Press, 1999); Ronald N. Giere, "The Role of Agency in Distributed Cognitive Systems," *Philosophy of Science* 73, no. 5 (December 2006): 710–19.

This approach, I think, helps to support my argument for the similarities, or even the fundamental sameness, of individual and group level cognition, mind, and agency. Opponents of collective mentality, group minds, and group agents have often pointed to a supposedly evident difference between individual and collective minds and intentions: individual human beings have a unified, centralized subjectivity and point of view. But research indicating the extent to which individual human cognition is in fact distributed, even within the limits of the neural system, challenges this long-standing assumption. According to Clark:

> Humans belong to the interesting class of what I'd like to call open ended ecological controllers. These are systems that seem to be specifically designed so as to constantly search for opportunities to make the most of body and world, checking for what is available, and then (at various timescales and with varying degrees of difficulty) integrating it deeply, creating whole new unified systems of distributed problem-solving.[49]

We are already, it would seem, "natural-born cyborgs,"[50] constituted by and distributed across a system that is not entirely biological. Here we can see how a group agent, constituted by other minds and artifacts, could function as a single mental unit. When we reflect upon ourselves as conscious agents, it is tempting to think that what makes us an individual is somehow located within the confines of our brain or body, and that it remains stable and constant over time, but the HEM helps explain how we recruit the environment, and other minds, to extend our cognitive capacity.

In a group, as in an individual mind, there are numerous simultaneous processes operating in parallel to maintain the functioning of the group as a subject with a unified point of view. This is known as "parallel processing," and it is necessary for complex skills like

[49] Andy Clark, "Soft Selves and Ecological Control," in *Distributed Cognition and the Will: Individual Volition and Social Context*, ed. Don Ross et al. (Cambridge, MA: MIT Press, 2007), 103.

[50] Andy Clark, *Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence* (New York: Oxford University Press, 2003).

reading or playing basketball,[51] where many highly integrated systems need to function independently toward a single goal, in order to coordinate movements and actions. We may be able to attend to some, though not all, of the automatic processing we are doing; I may be focused on my hand reaching for my coffee cup or I may do it thoughtlessly, with my mind directed elsewhere. But that isn't to say that action is unconscious, or fully subconscious; it is a deliberate, directed action by myself as an agent to fulfil a desire that is not at the forefront of my stream of consciousness, but neither is it fully absent. At any given moment, "it is assumed that we can only attend to one thing at a time, but we may be able to process many things at a time so long as no more than one requires attention."[52] Our attention shifts rapidly, allowing for the appearance of multi-tasking, and emotions, thoughts, memories, and perceptions flit in and out of our stream of consciousness, unbidden and difficult to pinpoint, either by their arrival or departure.

Although we can be aware of and direct our attention deliberately, we simply are not doing that most of the time. It isn't just our autonomic functions like the beating of our heart that our individual selves, our conscious minds, that happen to be distributed among various control centers, "but all *kinds* of human activities turn out to be partly supported by quasi-independent non-conscious subsystems."[53] There is no one neurological location or subsystem that has the final say, and the coordination of actions and resources is mostly "decided" without any conscious thought on our behalf, though they are often determined or influenced by standing beliefs, desires, and memories. It makes intuitive sense to believe that "my body has a mind of its own,"[54] highlighting the tension between the idea that the human

[51] David LaBerg and S. Jay Samuels, "Toward a Theory of Automatic Information Processing in Reading," *Cognitive Psychology* 6 (1974): 295.

[52] LaBerg and Samuels, 295.

[53] Clark, "Soft Selves and Ecological Control," 110. Italics in original.

[54] Daniel Dennett, "My Body Has a Mind of Its Own," in *Distributed Cognition and the Will: Individual Volition and Social Context*, ed. Don Ross et al. (Cambridge, MA: MIT Press, 2007), 93–100.

brain is the sole center of mind and agency (which is used to deny group agency), when phenomenally we can recognize how little of our own action feels directly within our control, and we cannot find any sort of distinct location or system that could be the seat of such control, besides the assemblage of parts (artificial, biological, and social) that we happen to identify with and have deeply integrated at the moment.

What we think of as our unified consciousness and sense of agency is potentially the result of a collection of subsystems that are in fact themselves agential, and may even be "like" something. Recalling the quote at the beginning of the introduction, our own minds seem to be a collection of agents, whose experience we cannot access, and our own mental experiences would presumably be inaccessible to the agential parts. If we don't insist on the existence of a fully unified, central controller and ultimate authority within our individual bodies, it becomes much easier to imagine a group as a similar sort of dispersed, ecologically controlled decision-making system, a "soft self."[55]

A group agent, I think, is a similarly soft self, with distributed processing systems instead of a system with distinct central controller. Even in the most centralized, hierarchical group, not every intention or mental state belongs to a single individual; various tasks, systems, and processes that direct the group and function to create the beliefs and desires of the group are distributed across the members. And they may be distributed in such a way that the group subject possesses a collective belief that is not held by a single member. Remember Gilbert's disjunctive criterion, where "a joint commitment to believe that p as a body does not require each participant personally to believe that p."[56] It is possible for two distinct groups with opposing beliefs to be composed of the exact same individuals.[57] Not only is this a

[55] Clark, "Soft Selves and Ecological Control."
[56] Margaret Gilbert, "Collective Epistemology," *Episteme* 1, no. 2 (October 2004): 101.
[57] Gilbert, 98.

useful criteria for resisting reductionism, but I think it interestingly parallels the distribution of mental states among our various individual subsystems.

So far I have covered that if the PP holds, the mind is socially extended. And if the mind is socially extended, there are group mental states. In this section I addressed the specifics of how a group mental state might function, because opponents of truly collective intentional mental states might claim that they are too different from the unified perspective of our individual mental states. But I have argued that, in fact, our own mental states are similarly distributed among various agential, or semi-agential parts, that nonetheless give rise to an apparently unified phenomenal experience. We should not assert the non-existence of a group mind, when it seems that even our own minds function as an integrated group of subsystems. And I do not think that group mental states are restricted to just those of knowledge and non-phenomenal experiences, so in the next section I will address the potential for groups to have more complex mental states.

Genuinely Collective Group-Level Emotions and Beliefs

In order to provide more thorough explanation of the social extension of mental states necessary for a robust, minded group agent, I will explore some theoretical and psychological evidence for genuinely collective mental states in this section. Because, although distributed information processing and memory retrieval systems are very interesting aspects of group cognitive systems, they are lacking the affective or phenomenal characteristics that we generally think intentional subjects possess. And, looking forward to the next chapter where I consider the challenge that consciousness poses to accounts of group agency and the HEM, I would like to lay some groundwork for genuinely collective mental states that might feel "like" something to the group agent, and are at the very least irreducible to the individuals of the group.

Psychological research has given some attention to the topic of collective emotion, and tentative conclusions suggest that group level emotions are: distinct from those of component individuals, shared socially, depend on level of group identification, and serve to direct and regulate group actions and attitudes.[58] Collective emotions are understood to exist on a spectrum from weakly to strongly collective, and this can be explained by background cognitive and affective capacities for collective emotions that function to motivate and justify collective actions.[59] According to this view, these collective emotions may precede and help give rise to collective intentions. This seems is compatible with Schmid's theory of plural pre-reflective self-awareness, and recalling how babies appear to share and experience emotional and affective states before they even perceive themselves as agents, arguably before they even qualify as conscious. Perhaps collective emotional experience is pre-conscious, subconscious, or unconscious, but it certainly plays an important role in consciousness.

One doesn't necessarily need to support the idea of a collective mind or consciousness in order to acknowledge the reality of genuinely collective group emotions, because likely not all emotional states are conscious. For example, Bryce Huebner proposes that groups possess emotions equivalent to the sort of unconscious, sub-personal emotions and moods of individuals that serve a particular sort of function; whether or not we are aware of them at the level of stream of consciousness, "they produce action-oriented representations that

[58] Eliot R. Smith, Charles R. Seger, and Diane M. Mackie, "Can Emotions Be Truly Group Level? Evidence Regarding Four Conceptual Criteria," *Journal of Personality and Social Psychology* 93, no. 3 (September 2007): 431–46; Diane M. Mackie, Lisa A. Silver, and Eliot R. Smith, "Intergroup Emotions: Emotion as an Intergroup Phenomenon," in *The Social Life of Emotions*, ed. Larissa Z. Tiedens and Colin Wayne Leach (Cambridge: Cambridge University Press, 2004).

[59] Mikko Salmela and Michiru Nagatsu, "Collective Emotions and Joint Action," *Journal of Social Ontology* 2, no. 1 (2016): 33–57.

reflexively reorient behavior in ways that facilitate coping with emotionally salient stimuli."[60] And while he doesn't think that collective phenomenal consciousness is a logical impossibility, he does not think that any existing groups are a suitable substrates, because they do not have sufficient bandwidth or integration.[61]

The capacity for collective emotional states might be necessary for a conscious subject, but it certainly isn't sufficient. The broader point that I am trying to make here is that groups, at the very least, meet not only the minimal conditions for mind and agency, but actually exhibit many more complex mental states that we think only belong to conscious subjects. It may be as Huebner claims that groups do have sufficient bandwidth, complexity, and informational integration to hold genuinely collective emotional states, though not enough to qualify as conscious.[62] I question on what basis philosophers make these claims regarding the quantitative aspects of different mental states, and consciousness in particular, and I suspect that more work needs to be done in this area to make this sort of definitive claim. If we are going to rule out group minds on such grounds, I think that we first need to be clear about how interpersonal information processing is measured, and what level is necessary for subconscious, and conscious experience.

Another complex, standing mental state that we consider part of the experience of a mind, subject, or agent is the ability to hold beliefs. So, can groups hold genuinely collective beliefs? At least on an interpretationalist view of collective intentions, where the goal is just to successfully interpret the behavior of an agent, group agents are arguably "true believers."[63] But in this case I will continue to take the functionalist position, due to its

[60] Bryce Huebner, "Genuinely Collective Emotions," *European Journal for Philosophy of Science* 1, no. 1 (January 1, 2011): 106.

[61] Huebner, 103.

[62] Huebner, 89.

[63] Deborah Tollefsen, "Organizations as True Believers," *Journal of Social Philosophy* 33, no. 3 (Fall 2002): 395–410.

popular acceptance and widespread use in the sciences, and think the that understanding of a belief makes sense in terms of how those beliefs play the proper role in the actions of group agents.

From a functionalist point of view, group agents with collective beliefs strongly implies a collective mind, the reductionist and plural subject accounts of group agency tend to deny the existence of group beliefs,[64] or somewhat radically depart from our common understanding of what belief is, by proposing conditions for group belief based on a background of collective acceptance.[65] And if group belief is analogous to individual belief, then a commitment (or acceptance) must not be necessary for belief.[66] It seems to me that if we speak of group mental states like emotions and beliefs, they should be analogous to individual states, otherwise we are misusing the terminology. If we do intend to use them analogously we should be willing to ascribe beliefs, acceptances, and actions to groups, not just as plural subjects, but as "supraindividuals,"[67] with a unified, rational perspective, and perhaps even phenomenal experience. There are other ways to refute this conclusion, like to claim that mental representations require neural mechanisms (i.e., brains), so even if mental representations are possible via different substrates, they would be nothing like human representations, emotions, beliefs, and desires.[68] These claims, however, are not uncontroversial, and maybe be opposed by either rejecting that beliefs require representations, or asserting the existence of (or possibility for) collective representations.

Our own seemingly well-integrated representational mental states are the product of many lower-level representational and quasi-perceptual states, so collective mental states

[64] Bratman, "Shared Intention."

[65] Gilbert, *Joint Commitment*.

[66] Frederick Schmitt, "Group Belief and Acceptance," in *From Individual to Collective Intentionality* (Oxford: Oxford University Press, 2014), 61–96.

[67] Schmitt, 88.

[68] Paul Thagard, *The Brain and the Meaning of Life* (Princeton, NJ: Princeton University Press, 2010), 275.

should not be ruled out prima facie.[69] HEM and distributed cognition provide one promising avenue for envisioning collective representations, and some empirical work is being done toward this end.[70] A robust notion of collective intentions, emotions, beliefs, and representations in turn paves the way for a realist conception of the group agent.

The Irreducible, Non-Emergent Group Agent

Mental states are all, or at least primarily, considered to be intentional states. Using the PP, we are able to recognize the social extension of mental states, and I think this supports the idea that a group agent is the subject of these collective mental states. I previously addressed some challenges to reductive theories of group agency, and here I would like to provide one further theory to more strongly establish the irreducible nature of the group agent, as a robust subject of experience, with mental states and minds of their own.

The idea that the mind extends beyond the brain has many implications for our concepts of mind and agency, particularly when we understand that the mind can extend socially, and that this process mirrors the distributive functioning of our own minds. The traditional individualist approaches to the HEM and group agency tend to reject the possibility that their theories imply a collective consciousness, but in this section I will argue for a strong notion of a group agent, and that we should not reject the possibility of it as even a conscious entity.

Typically, realist accounts of group agents portray them as something "over and above" their individual members, that emerges from the constituent parts. In this section I would like to suggest that this sort of metaphysical relationship isn't necessary for a robust

[69] Bryce Huebner, "Do You See What We See? An Investigation of an Argument Against Collective Representation," *Philosophical Psychology* 21, no. 1 (February 1, 2008): 91–112.

[70] Jeroen de Ridder, "Representations and Robustly Collective Attitudes," in *Socially Extended Epistemology*, ed. J. Adam Carter et al. (Oxford University Press, 2018); Tibor Bosse et al., "Collective Representational Content for Shared Extended Mind," *Cognitive Systems Research* 7, no. 2 (June 1, 2006): 151–74.

group agent. When people claim that group agents have emergent characteristics, this can be understood in one of two ways. They might mean that the higher-level qualities are novel insofar as they are not found at the lower level, and although they should be analyzed at the group-level, a sufficiently advanced modeling system would allow for accurate prediction; this is often referred to as weak emergence. Or they may view the characteristic of group agents as strongly emergent, meaning that they could never be predicted, even with complete knowledge of the individual members.

Both weak and strong emergentists usually claim that their theories are non-reductive and non-aggregative, but weak emergentism can be challenged on this, because if a system can be predicted based on the smaller (or smallest) components, then it seems their position is compatible with certain reductionist views. It may be that only strong emergence is the only truly emergentist position. Both variations assume a metaphysical hierarchy with different levels of organization. The weak emergentist relies on the microphysicalist assumption that causal powers lie at the lowest level of organization and the strong emergentist can be accused of incompatibility with science, since they seem to be asserting the complete unpredictability of certain characteristics such that they take on a supernatural quality. An emergentist view of the group agent must either posit that the actions and intentions of the agent can be predicted by the characteristics of the individual members, meaning that they are not truly novel, and thus collapse into a refined reductionist position, or assert that the agency of the group cannot be explained by those of the individuals, relegating the relationship between the parts and the whole to eternal mystery.

One objection to group agency is that there is no need, or explanatory role, for a group agent "over and above" the individual because all of the intentions and actions of the group

are analyzable at the individual level.[71] This sort methodological individualism is the sociological equivalent of a microphysicalist viewpoint, asserting that all causal power and action resides within the smallest proper parts. But I do not feel compelled to take these metaphysical views for granted. Mariam Thalos proposes a flat ontology, a "scale-free universe model," that rejects the notion of different "levels," which are the foundation of emergentist theory.[72] I am inclined toward this view because it is consistent with both materialism and modern physics, or so Thalos argues, quite convincingly. According to her hypothesis, activity and causation "happens at every scale of measurement, every scale of organization in the universe,"[73] meaning that there is no ultimate, fundamental level or privileged scale. She argues that the commonly accepted hierarchical metaphysical model of the universe has been a burden to both philosophy and physics, creating imaginary boundaries and blinding us to the interactions between the supposedly distinct levels.

Relying on this ontological view of action within the universe, I would like to suggest that we cannot claim the agency, intention, or consciousness of the individual organism is the fundamental level where those processes occur. We should not assume that action happens at only one particular level of organization, or in only one direction, and to do so seriously limits our potential ability to understand the complex nature of interactions between entities. Assuming that a group of people function as an agent, there is no good reason to deny agency at any particular level based on nothing other than its relative scale to other presumed agents. When we reimagine our ontology in this way, there is no need to posit that the group agent is

[71] Kirk Ludwig, "Is Distributed Cognition Group Level Cognition?," *Journal of Social Ontology* 1, no. 2 (January 1, 2015); Sara Rachel Chant, "The Special Composition Question in Action," *Pacific Philosophical Quarterly* 87 (2006): 422–41.

[72] Mariam Thalos, *Without Hierarchy: The Scale Freedom of the Universe* (Oxford: Oxford University Press, 2013).

[73] Thalos, 5.

anything "over and above" the individuals, and thus avoiding the anti-emergentist critiques of group agency.

With this non-hierarchical ontology in mind, I would like to argue for a robust group agent that functions and behaves as an agent does, with the requisite mental states, beliefs, desires, and emotions that provide reasons and determine behavior, that operates with a rational point of view, and may even have some sort of inner phenomenal experience. If the functionalist is willing to accept Clark and Chalmers's PP for cognition, we should also allow for instances of extended agency on the same terms that we do for the extended mind and cognition. We don't have first-person access to the internal life and sense of agency of other individuals, or the group agents that we constitute. So, I think there is no other reasonable way to ascribe agency to another system other than it operates as an agent, and that this is important for both ethical and practical reasons. And in many cases, we can show that groups do exhibit these qualities of agency to various degrees, and sometimes a very high one. Remember Pettit's functionalist definition of the intentional subject as a kind of agent that acts according to their beliefs and desires in a way that makes sense with their values, and groups can easily meet these stipulations. Thus far, I hope I have shown that literally ascribing mental states to groups is not as absurd as it is often portrayed.

One remaining challenge to the notion of a robust group agent is the question of where we should draw the boundary of the individual agent and the mind, if they extend in the way that has been suggested. If the mind extends into the environment, and an agent can extend into many people, where does it stop? Perhaps these conceptions of mind and agency are overly liberal, to the point that they no longer pick out the interesting phenomena that we set out to understand and explain. Although this subject could use a great deal more exploration, there are some promising suggestions for the mechanisms by which the agent is constructed and defined, internally and externally. For example, we might look to accounts of

the narrative construction of the self at both the level of the individual and the group,[74] self-by-doing,[75] personal identity and the ethical criteria of personhood,[76] or rational unity.[77] If our own sense of identity and agency is any indication, we do extend into the world and even the people around us, but not infinitely, and I can imagine that there are many different mechanisms and processes that contribute the unification of a group agent.

[74] Shaun Gallagher and Deborah Tollefsen, "Advancing the 'We' Through Narrative," *Topoi* 38, no. 1 (March 1, 2019): 211–19; Deborah Tollefsen and Shaun Gallagher, "We-Narratives and the Stability and Depth of Shared Agency," *Philosophy of the Social Sciences* 47, no. 2 (March 2017): 95–110; Catriona Mackenzie, "Embodied Agents, Narrative Selves," *Philosophical Explorations* 17, no. 2 (June 2014): 154–71; Richard Heersmink, "The Narrative Self, Distributed Memory, and Evocative Objects," *Philosophical Studies* 175, no. 8 (August 2018): 1829–49.

[75] Stephan Verschoor and Bernhard Hommel, "Self-By-Doing: The Role of Action for Self-Acquisition," *Social Cognition*, December 23, 2017.

[76] Carol Rovane, *The Bounds of Agency: An Essay in Revisionary Metaphysics* (Princeton, NJ: Princeton University Press, 1998).

[77] Pettit, "Groups With Minds of Their Own"; Philip Pettit, "The Reality of Group Agents," in *Philosophy of the Social Sciences*, ed. C. Mantzavinos (Cambridge: Cambridge University Press, 2009), 67–91.

Chapter Three: The Challenge of the Collective Mind

The Question of Consciousness

Thus far, I have established that if we accept the PP, as I think we should, then the mind can extend beyond the individual biological organism and into the social world. The social extension of the mind allows for the existence of group mental states, which I argue are necessary for, and provide evidence for a robust notion of the group agent. However, opponents of group agency and the HEM both argue that there cannot be talk of group cognitive or mental states without implying group minds or consciousness. For example, Richard Rupert points out that the most commonly accepted characteristics of mentality and cognition to this point have been the experience of phenomenal consciousness, flexible intelligence, and representational capabilities, which are essential properties of minds that groups just don't possess.[78] He also says that "No group mind claims to have conscious experience,"[79] but I wonder if they don't, and how they would communicate such an experience to us.

A great many of the challenges posed to the HEM (and group agency), I think, originate from the desire to reject the notion of an extended mind or agent because that implausible conclusion undermines those theories, and make our notions of mind and agency overly liberal, or perhaps meaningless. It does seem a bit ludicrous to suggest a genuinely collective mental state, group mind, or group agent, because we don't recognize those features in a group in the same way that we do ourselves and other people. We may say that a group has a feeling, but we don't see it in their facial expression; we communicate with groups, but they do not tell us they are conscious. Critics may accept the PP, but deny that that group interactions function in the same way our brains do, allowing our minds to exhibit

[78] Robert Rupert, "Minding One's Cognitive Systems: When Does a Group of Minds Constitute a Single Cognitive Unit?," *Episteme* 1, no. 3 (February 2005): 177–88.

[79] Rupert, 178.

true intentionality, mind and consciousness in a way that would be impossible for a collective. But in the previous chapter, I hope that I provided some evidence that our minds function in a distributed, parallel manner quite similar to the cognitive operations of a group.

In this chapter, I will address the primary objection to the idea of a robust, thinking, group agent, which is that discussion of group mental states and agency insinuates the existence of an impossible collective consciousness. I don't wish to object to that conclusion inasmuch as I wish to resist the assumption that it is inherently impossible. This objection is also leveraged against the HEM independently, which is interesting. If two explanatorily powerful theories both imply the possibility of a genuinely collective mind, I think that it is a possibility that should be taken seriously instead of dismissed offhand, as it usually is. Even theorists whose work does strongly indicate a group subject or collective experience are quick to distance themselves from the idea.

Clark and Chalmers think that even though cognition extends, that consciousness decidedly does not. Dennett confidently proclaims that it isn't like anything to be a colony of ants or a brace of oxen, though it may be like something for the individual creatures.[80] Searle says that discussion of the group minds is "at best mysterious and at worst incoherent," and places advocates of theories of collective consciousness in opposition to "empirically minded" philosophers.[81] Although not an advocate for the Group Mind Hypothesis (GMH) himself, Huebner declares that:

> No materialist account of consciousness can rule out the *psychological possibility* of collective consciousness a priori; and there is reason to believe that the most promising materialist theories of consciousness imply the *possibility* of collective mentality even in the actual world.[82]

[80] Daniel Dennett, "Annual Question: What Do You Believe Is True Even Though You Cannot yet Prove It?," Edge.org, 2005, https://www.edge.org/response-detail/11902.

[81] Searle, *Consciousness and Language*, 93.

[82] Bryce Huebner, *Macrocognition: A Theory of Distributed Minds and Collective Intentionality* (Oxford University Press, 2013), 118. Italics in original.

Considering the prevalence of the materialist view of consciousness, this possibility seems like something to consider, at the very least. Philosophers have been happy to consider the extension of the mind and cognition, group-level intentions, agency, memories, and emotions, and even "groups with minds of their own,"[83] but they have overwhelmingly drawn the line at similarly extending consciousness beyond the individual brain. Although I will not assert the definite existence of any form of collective consciousness, I will address the most prevalent arguments for rejecting the possibility, to see whether they are justified.

One notable exception to the assertion that consciousness occurs in the brain is Alva Noë, who challenges the foundational assumption of mainstream neuroscientific psychological disciplines and the philosophy of mind. Noë thinks that "we have been looking for consciousness where it isn't. We should look for it where it is. Consciousness is not something that happens inside us. It is something we do or make. Better: it is something we achieve. Consciousness is more like dancing that it is like digestion."[84] Looking for consciousness in the brain has been an unsuccessful endeavor, because it isn't there. Instead, consciousness just is the dynamic relationship and interaction of our brains, bodies, and the environment, where these pieces are all necessary and inextricable. Drawing on a great deal of empirical work, Noë makes a powerful case for a paradigm shift in how we view and study the mind and consciousness. According to this view, not only are the environment and other people occasionally coopted and integrated into our cognitive processing (as in HEM), but they are absolutely essential to its constitution and function. If we think through the brain-in-a-vat thought experiment we recognize that it isn't an isolated mind or disconnected consciousness, but that the vat itself would effectively need to be a living body situated in an

[83] Pettit, "Groups With Minds of Their Own."

[84] Alva Noë, *Out of Our Heads: Why You Are Not Your Brain, and Other Lessons from the Biology of Consciousness* (New York: Hill and Wang, 2009), xxi.

environment;[85] there is not untethered, autonomous Cartesian mind, and it does not reside strictly within the tissue of our brain and nervous system. For this reason, I think we should also be skeptical that it only resides inside the individual, as the boundaries of the brain and body are somewhat arbitrary.

Contrary to Searle's claim that discussion of a collective mind is unscientific, it would appear that some empirical research does point to the incredible plasticity regarding what we perceive to be part of our bodies, minds, and sense of self; these borders are constantly shifting and allowing us to adopt and modify parts of the environment, and to extend our thought processes and into the world, physically and socially. Noë also touches upon the relationship between parents and newborns as a particularly unified dyad, where a great deal of the infant's functioning (physiologically, mentally, and emotionally) does not occur within the infant themselves, but instead in this highly sensitive and dynamic relationship in which the baby is eventually able to recognize their individuality, or more accurately, become more well-integrated into their environment.[86] I previously mentioned infant development because the close relationships between babies and caregivers are highly intimate, entwined, and responsive, so that they seem to exemplify a socially extended mind, and therefore a collective mental state.

I also discussed the possibility of socially extending the mind and cognition, but with Noë in mind, it seems that maybe everyone has prematurely rejected the extension of consciousness. Perhaps consciousness and agency alike come in degrees. I will focus on this question of phenomenal consciousness of the group in the final section of this chapter, but first I would like to address the characteristics that we normally consider to be significant

[85] Noë, 12–13.
[86] Noë, 50–51.

about consciousness, and to point out how groups may possess at least some of them in a genuinely collective sense.

Anti-Nesting Principles

One objection to the idea of a group mind, or consciousness, are "anti-nesting principles,"[87] which assert that one consciousness cannot be composed of, or nested within, another conscious organism. This is articulated briefly by both Hilary Putnam and Guilio Tononi, who argue that consciousness can only exist at a single level of organization.[88] Eric Schwitzgebel addresses this objection in his paper "If Materialism is True the United States is Probably Conscious," where he provides a challenge to materialism, explaining that if "we set aside our morphological prejudices against spatially distributed group entities, we can see that the United States has all the types of properties that materialists tend to regard as characteristic of conscious beings."[89] And this is a conclusion that functionalists and materialists would rather avoid, because, as I have pointed out, collective consciousness simply sounds too far-out and unscientific, if not downright impossible. Although Schwitzgebel's paper is intended as a critique of materialism, I think that it actually provides us with some excellent examples of how the mind could be extended and distributed spatially, while allowing for conscious experience that superficially looks much different from our own.

The two examples he uses of conceivable alien beings to make his point are Sirian supersquids and Antarean antheads. Instead of a centralized brain, supersquids have nodes

[87] Eric Schwitzgebel, "If Materialism Is True, the United States Is Probably Conscious," *Philosophical Studies* 172, no. 7 (2015): 1702.

[88] Hilary Putnam, "Psychological Predicates," in *Art, Mind, and Religion,* ed. W. H. Capitan and D. D. Merrill (Liverpool: University of Pittsburgh Press, 1965); Giulio Tononi, "The Integrated Information Theory of Consciousness: An Updated Account," *Archives Italiennes de Biologie* 150 (2012): 290–326.

[89] Schwitzgebel, "If Materialism Is True, the United States Is Probably Conscious," 1698.

distributed among thousands of detachable tentacles that can communicate directly with one another using light signals. However spatially dispersed they are, in their communications with us, they report having a fully integrated, unified stream of experience. The Antareans resemble wooly mammoths, but their cognitive processes are much slower than our own, and instead of neurons, their heads are full of tiny insects with their own individual sensory experiences, whose complex interactions are functionally equivalent to our own brains. As he points out, there is nothing inherent to a materialist or functionalist view of consciousness that would rule out the plausibility of multiple levels of consciousness. For the same reason, I think the materialist or functionalist should also admit that group minds and agents are at the very least a possibility.

For this reason, it seems to be difficult to take the anti-nesting challenge very seriously, because no argument has been provided for why a larger consciousness could not be composed of smaller, conscious parts. And I don't think we even need to look so far as to fantastical alien species to see how one seemingly unified organism is composed, in fact, of many different organisms. For example, there are approximately an equal number of human cells and bacterial cells in the human body,[90] and those microbial cells are absolutely essential to our survival. And many other examples of symbiotic organisms exist, from lichen to the gut flora of ruminants. That is not to say that any of these microorganisms are conscious, but if we can be composed of smaller life forms whose functions we are unaware of, that play necessary roles in our proper functioning, I see no reason to think that those smaller parts could not themselves be conscious just because they reside inside a larger organism. Earth itself, from a distance, might be seen as a sort of super-organism, of which we are only a small part. There are other ways to deny that our constituent parts are conscious, like that

[90] Ron Sender, Shai Fuchs, and Ron Milo, "Revised Estimates for the Number of Human and Bacteria Cells in the Body," *PLoS Biology* 14, no. 8 (August 19, 2016).

they lack sufficient information processing and integration capacities. But I do not think consciousness can be ruled out just because of the location inside of another consciousness, any more than it can be ruled out because it does not have a brain like ours.

I think it is important for us to reflect upon our anthropocentric assumptions about consciousness, and to resist the temptation to believe that consciousness in other beings must look very much like ours in terms of size, location, spatial arrangement, processing speed, or composition. I can accept that there is a certain necessary level of integration necessary for consciousness, but I see no convincing argument for why consciousness could not exist at a very small, or incredibly large scale, be spatially distributed, slow, or consist of parts that are themselves conscious. We really don't know what characteristics about the human brain, body, and interaction with the world give rise to our apparently unique sort of consciousness and experience of the world, so it seems very preemptive to deny the possibility of consciousness to other kinds of entities, like groups, particularly when there is an argument to be made that groups are capable of acting as agents, and even as minds.

Sense of Self and Agency

Even proponents of a non-minimal accounts of group minds, who argue that groups can and do display the characteristics of individuated minds (epistemologically distinct, unified, and integrated rational points of view) make claims like "there is obviously no first-personal self-hood in groups,"[91] and deny that groups are capable of conscious experience or phenomenal integration. This claim, however, is again not so much argued as it is asserted. And because of this, I continue to wonder about the supposedly obvious truth of the claim.

This also raises another important question regarding how much, if any real "self" is necessary for conscious experience or agency. It seems reasonable that a group agent could

[91] Thomas Szanto, "How to Share a Mind: Reconsidering the Group Mind Thesis," *Phenomenology and the Cognitive Sciences* 13, no. 1 (March 2014): 107.

exist in a robust sense, even as a conscious being of a sort, that may not have a strong sense of self. Because even at the individual level, "just because the subjective character of consciousness gives rise to a *sense* of self—that is, the felt sense of being a stable who, or owner of conscious episodes—it does not follow that this who really exists in any autonomous or enduring sense."[92] Instead of viewing consciousness as something that belongs to the self, perhaps we can change perspective and see the self as a result of a certain kind consciousness, and avoid over-emphasizing its metaphysical significance.

Luke Roelofs similarly defends the possibility of "moderately selfless agents" that exhibit a form of "seamless" agency, similar to our individual agency in that "it does not involve any awareness of being based in the agency of distinct, self-aware, parts or members of us."[93] Note that this doesn't assert that we are not composed of different agential parts ourselves, just that our awareness does not extend to them. Just as the experience of a group agent need not include or be accessible to its components, or vice versa. Although you and I each have a distinct sense of our own agency, that does not rule out the prospect of viewing ourselves as composite agents. In fact, I think that is likely closer to the truth, which I elaborated on in the second chapter.

It is not only our sense of self, but even our own sense of agency that can be challenged. According to research done by Thomas Metzinger, approximately two-thirds of our waking life is spent mind-wandering, and when we are in that state we don't exhibit the qualities normally attributed to mentally autonomous cognitive systems, like "mental agency, explicit, consciously experienced goal-directedness, or availability for veto control," leading

[92] Joel Krueger, "The Who and How of Experience," in *Self, No Self? Perspectives from Analytical, Phenomenological, and Indian Traditions*, ed. Mark Siderits, Evan Thompson, and Dan Zahavi (New York: Oxford University Press, 2011), 27.

[93] Luke Roelofs, "Rational Agency without Self-Knowledge: Could 'We' Replace 'I'?," *Dialectica* 71, no. 1 (2017): 5.

him to conclude that "personal level cognition is an exception rather than the rule."[94] This raises many questions for me regarding the standards that an organism must meet in order to qualify as a thinking, conscious, self-aware agent, considering that human beings, typically the paradigmatic example of such, often do not meet those standards ourselves.

Metzinger also conducts psychological studies on human dream states, noting the similarities between dreaming and mind-wandering. He attempts to identify the necessary conditions of "minimal phenomenal selfhood" that are necessary for the existence of a first-person perspective, or the most basic form of self-consciousness.[95] This is, predictably, an enormous undertaking that will take a great deal of work to properly understand. But there are some very important philosophical implications of research like this, and how we think about consciousness and agency. Dreams are a good example of how we can be very minimally self-aware. We are not totally unconscious when we are asleep, we have some sort of bodily awareness and control that is necessary for us to stop us from falling out of bed, and awareness enough to be woken up. The borders of consciousness are not clear and distinct. And I think that we should not apply those standards to other organisms, if they mean that a great deal of our own lives we would not be considered conscious agents.

Access and Phenomenal Consciousness

The final objection to the group mind and the robust group agent is the claim that regardless of the other characteristics of the group agent, it is clearly not phenomenally conscious. It is difficult to deny that groups possess informational content that is available to

[94] Thomas Metzinger, "The Myth of Cognitive Agency: Subpersonal Thinking as a Cyclically Recurring Loss of Mental Autonomy," *Frontiers in Psychology* 4 (December 19, 2013): 1.

[95] Thomas Metzinger, "Why Are Dreams Interesting for Philosophers? The Example of Minimal Phenomenal Selfhood, Plus an Agenda for Future Research," *Frontiers in Psychology* 4 (2013).

influence thoughts and actions, but this kind of "access consciousness" (A-consciousness) may be distinct from "phenomenal consciousness" (P-consciousness).[96] According to Block, P-consciousness is the experiential form of consciousness which is "like" something.[97] He argues that these two kinds of consciousness are separable, and suggests an everyday example of P-consciousness without A-consciousness where you don't actively notice the construction noises outside of your window, and suggests that A-consciousness without P-consciousness is "conceptually possible," although he does not have an actual example of such a mental state, he does want to argue that A- and P-consciousness are fully distinct. Some accounts of group agency accept this distinction and reject the possibility of group P-consciousness, but accept or advocate for A-conscious group mental states.[98] Those hesitant to resist the GMH or the possibility of a fully conscious group agent may claim that the contents of the extended mind, and social group-level mental states, exemplify A-consciousness, without implying P-consciousness. I would like to address this very reasonable suggestion in two steps.

First, I would like to address the concern over the apparently lacking expression of a phenomenally conscious experience on the part of group agents. Because, whether a group is just an aggregate or a complex and integrated group agent, we do communicate with them frequently. And it isn't always clear which sort of group we are interacting with. A group of friends might say "we're so happy for you," and just be referring to the aggregation of their individual joys, but one nation may impose sanctions on another, outwardly claiming some just cause, yet acting for their own benefit and longevity as a unified agent. A group has

[96] Ned Block, "On a Confusion about a Function of Consciousness," *Behavioral and Brain Sciences*, 1995, 227–87.

[97] Thomas Nagel, "What Is It Like to Be a Bat?," *The Philosophical Review* 83, no. 4 (October 1974): 435–50.

[98] Deborah Tollefsen, *Groups as Agents* (Cambridge, MA: Polity Press, 2015); Christian List, "What Is It Like to Be a Group Agent?," *Noûs* 52, no. 2 (June 2018): 295–319; Huebner, *Macrocognition*.

never come out and said, “I’m a phenomenally conscious agent,” but that is not typically how we ascribe agency to other people or animals anyhow. Our ascription of mind and consciousness to other organisms and other people does not seem to be very dependent on their outright declaration of their conscious experience, so I do not think we should hold groups to that standard either. I believe that my friends, neighbors, and pets have a phenomenal experience of the world, although they have never told me, but I do believe it is like something to be them based on their responses and interactions with their environment.

We may communicate with groups through spoken and written language, and while they may not explicitly claim a phenomenal experience, they do express many thoughts, beliefs, and even emotions that we would normally assume to be an indication of conscious experience if they were communicated by an individual person. Learning to identify group agents, particularly robust ones with a variety of integrated mental processes that operate as a unified agent over time, will certainly be a challenge. How could we attempt to verify the conscious experience, or even sense of agency, of a being of significantly different scale and composition than ourselves, particularly one that we are a part of? Because the GMH and conceptions of robust group agents have not been given much serious attention, a great deal more work could be done on how to recognize a large (or small) scale mind, agency, and consciousness. But there is some promising early research in this area, and I will return to this discussion in the conclusion.

Second, I wish to challenge Block’s assertion that A- and P-consciousness are fully distinct, and that one could exist without the other. It is not obvious to me how there could be truly group-level mental states without any degree of group-level phenomenal consciousness. If there is a non-reductive collective emotion or belief that is at all analogous to our own, to experience an emotion or to recall a belief feels like something. A memory or belief must belong to a subject who can refer to it and act on it in relevant ways, otherwise that sort of

wording is inaccurate. Even a computer may qualify as A-conscious, but we do not ascribe mental states or agency to a computer, or claim that they remember or feel emotion, as we do about groups on a daily basis. We "cannot explain what it is to consciously reactivate or access a belief content in terms of A-consciousness alone," therefore, there must be some P-conscious agent to access the belief, since "a state is not a belief unless the owner of the state is disposed to access the state's content in a corresponding conscious judgment. Thus, if there is no such thing as group consciousness, then we cannot literally ascribe beliefs to groups."[99] To claim that a group is strictly A-conscious with no level of P-consciousness seems to be sidestepping the problem that there still needs to be a P-conscious agent to do the accessing. I question whether a strictly A-conscious organism could exist, and if it did, how it could be considered conscious.

You may still insist that it is impossible for a group to have a similar phenomenal experience to your own, which is very vivid. But it is also very difficult to explain how my own conscious experience exists, and why it feels likes like something to be me, and how I can confidently claim that it doesn't feel like anything to be an amoeba, or a part of my brain. Our assertions and ascriptions of consciousness seem to be based on very little other than denying that it is possible for any organism that is much different from ourselves. I don't disagree that there is something very interesting, or possibly unique in our world, about human consciousness. I have a strong sense of what it is to be me, I deeply feel my emotional states, my sensory input, the positioning of my body, and so forth, and that seems to clearly establish my existence as a metaphysically distinct unit. But I don't think that feeling alone should dominate the debates around group cognition and agency. After all, as I established in the second chapter, it would seem that I am nothing more than a very well-integrated set of

[99] Søren Overgaard and Alessandro Salice, "Consciousness, Belief, and the Group Mind Hypothesis," *Synthese*, March 1, 2019.

subsystems myself, and each of those, or some collection of them, may feel like something in their own right.

Conclusion

When List asks, "what is it like to be a group agent?" and concludes that it isn't like much, if anything,[100] he does think that this is a question that can be analyzed scientifically. Referring to neuroscientific work on integrated information theory (ITT) that seeks to pinpoint the physical structures necessary for consciousness,[101] List claims that the level of informational integration within a group is too low to produce a significant degree of phenomenal consciousness. Although there are many strengths to ITT, like a graded scale that corresponds to higher and lower levels of consciousness that can explain seemingly borderline cases like experiences of general anesthesia and sleepwalking, measuring the integration of information within any system is difficult and imperfect. List claims that the measurement of integration within a group would be low (though not zero), but I have not found any discussion of how this data is collected or analyzed, and thus his conclusion seems like a premature speculation. How can we attempt to measure the amount of information simultaneously transferred between people, not only with language, but facial expressions, gestures, smell, touch, sound, base level similarities, shared background knowledge and assumptions?

However, I think it is very important that List does take the possibility of collective consciousness seriously, opening the door for other creative approaches to looking for it, and the understanding consciousness and agency on a graded scale. Instead of looking for "the mark of the mental" or the "mark of the cognitive" we might be better served looking for specific functions like attention, problem solving, and memory that systems may exhibit at different levels, but that can be rightfully be understood as integrated cognitive systems.[102]

[100] List, "What Is It Like to Be a Group Agent?"

[101] Giulio Tononi and Christof Koch, "Consciousness: Here, There and Everywhere?," *Philosophical Transactions: Biological Sciences* 370, no. 1668 (2015): 1–18.

[102] Theiner, Allen, and Goldstone, "Recognizing Group Cognition."

Theiner refers to psychological work on transactive memory systems (TMS)[103] to show how a group of people can function as memory system that cannot be reduced to the cognitive functions of the individuals because "the groups' adaptability and information processing capacities arise from interactions among the individuals and their environments."[104] Although he doesn't discuss the GMH, I think this approach provides us with a powerful tool for looking for something that might deserve to be called a group mind. Perhaps it doesn't feel like anything to be a TMS, but what if that system has significant overlap and integration with other related systems? I think we need to leave room for the possibility that other kinds of systems may have some experience that may be difficult for to us to recognize.

Robust group agents (and the HEM) both strongly imply that a collective mind is at least conceptually possible. Perhaps there are group agents that possess collective minds, or at least some markers of conscious experience, whose experience is incredibly different from our own, but that it does feel like something to be. Rather than denying the possibility of this preemptively, we should strategize about how to look for systems like this. And that this work is particularly important to do as part of larger projects like looking for evidence of alien intelligence, or in attempting to create an autonomous, conscious, artificial intelligence.

My central argument in this thesis is that the PP can be applied to show how mental states can be socially extended in such a way that they function to cause the actions of robust, irreducible, non-emergent group agents. These mental states are in fact quite similar to our own, in that they are distributed among many parts that are simultaneously and independently processing information. Some of these subsystems might be considered agents themselves, and may be the independent subjects of experience that we are unaware of. Talk of group

[103] Daniel M. Wegner, Toni Giuliano, and Paula T. Hertel, "Cognitive Interdependence in Close Relationships," in *Compatible and Incompatible Relationships*, ed. William Ickes (New York: Springer Publishing, 1985), 253–76.

[104] Theiner, Allen, and Goldstone, "Recognizing Group Cognition," 390.

mental states does imply a sort of group mind or consciousness, which is typically rejected and considered an undesirable consequence, but I suggest that it should be seriously considered through the use of arguments that consciousness is not restricted to the brain and that groups are the subjects of genuinely collective emotions, intentions, and beliefs. I think that when these are considered together, they indicate that it may be like something to be a group.

This is an important consideration, because of our relative lack of success so far in understanding and anticipating the behaviors of groups. Even in purportedly democratically organized groups, the group as a whole may hold contradictory individual beliefs, and take actions that most, or even all of the group members disagree with. Instead of looking at the psychology of the individual to understand why this happens, I think we should move beyond the methodologically individualistic approach and seek to identify the kind of interactions and processes that might create a unified, minimally self-aware, collective subject. Tools and analysis like ITT and TMS may also help us look for, or create, consciousness or mentality in other places where we have assumed that it does not exist, like nested systems or systems of a much different scale from ourselves. Recognizing that even individual human agents are effectively "a society of different minds,"[105] I think that may help us recognize and accept the possibility of robust group agents composed of individual human minds.

[105] Minsky, *The Society of Mind*, 290.

References

Block, Ned. “On a Confusion about a Function of Consciousness.” *Behavioral and Brain Sciences*, 1995, 227–87.

Bosse, Tibor, Catholijn M. Jonker, Martijn C. Schut, and Jan Treur. “Collective Representational Content for Shared Extended Mind.” *Cognitive Systems Research* 7, no. 2 (June 1, 2006): 151–74.

Bouchard, Frédéric, and Philippe Huneman, eds. *From Groups to Individuals: Evolution and Emerging Individuality*. Vienna Series in Theoretical Biology. Cambridge, MA: The MIT Press, 2013.

Bratman, Michael. *Shared Agency: A Planning Theory of Acting Together*. New York: Oxford University Press, 2014.

———. “Shared Intention.” *Ethics* 104, no. 1 (1993): 97–113.

Chant, Sara Rachel. “Collective Action and Agency.” In *The Routledge Handbook of Collective Intentionality*, edited by Marija Jankovic and Kirk Ludwig, 13–24. New York: Routledge, 2018.

———. “The Special Composition Question in Action.” *Pacific Philosophical Quarterly* 87 (2006): 422–41.

Clark, Andy. *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA: MIT Press, 2001.

———. *Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence*. New York: Oxford University Press, 2003.

———. “Soft Selves and Ecological Control.” In *Distributed Cognition and the Will: Individual Volition and Social Context*, edited by Don Ross, David Spurrett, Harold Kincaid, and G. Lynn Stephens, 101–22. Cambridge, MA: MIT Press, 2007.

———. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Philosophy of Mind. Oxford: Oxford University Press, 2008.

Clark, Andy, and David Chalmers. “The Extended Mind.” *Analysis* 58, no. 1 (1998): 7–19.

Crane, Tim. “Intentionality as the Mark of the Mental.” In *Contemporary Issues in the Philosophy of Mind*, edited by Anthony O’Hear, 1998.

Dennett, Daniel. “Annual Question: What Do You Believe Is True Even Though You Cannot yet Prove It?” Edge.org, 2005. https://www.edge.org/response-detail/11902.

———. “My Body Has a Mind of Its Own.” In *Distributed Cognition and the Will: Individual Volition and Social Context*, edited by Don Ross, David Spurrett, Harold Kincaid, and G. Lynn Stephens, 93–100. Cambridge, MA: MIT Press, 2007.

Feldman, Ruth, Ilanit Gordon, Moran Influs, Tamar Gutbir, and Richard P Ebstein. “Parental Oxytocin and Early Caregiving Jointly Shape Children’s Oxytocin Response and Social Reciprocity.” *Neuropsychopharmacology* 38, no. 7 (June 2013): 1154–62.

Feldman, Ruth, Romi Magori-Cohen, Giora Galili, Magi Singer, and Yoram Louzoun. “Mother and Infant Coordinate Heart Rhythms through Episodes of Interaction Synchrony.” *Infant Behavior and Development* 34, no. 4 (December 2011): 569–77.

Gallagher, Shaun. “Direct Perception in the Intersubjective Context.” *Consciousness and Cognition*, Social Cognition, Emotion, and Self-Consciousness, 17, no. 2 (June 1, 2008): 535–43.

———. “The Socially Extended Mind.” *Cognitive Systems Research* 25–26 (December 2013): 4–12.

Gallagher, Shaun, and Anthony Crisafi. “Mental Institutions.” *Topoi* 28, no. 1 (March 2009): 45–51.

Gallagher, Shaun, and Deborah Tollefsen. “Advancing the ‘We’ Through Narrative.” *Topoi* 38, no. 1 (March 1, 2019): 211–19.

Gertler, Brie. “Overextending the Mind?” In *Arguing About the Mind*, edited by Brie Gertler and Lawrence Shapiro, 192–206. New York: Routledge, 2007.
Giere, Ronald N. “The Role of Agency in Distributed Cognitive Systems.” *Philosophy of Science* 73, no. 5 (December 2006): 710–19.
Gilbert, Margaret. “Collective Epistemology.” *Episteme* 1, no. 2 (October 2004): 95–107.
———. *Joint Commitment: How We Make the Social World*. New York: Oxford University Press, 2014.
———. *On Social Facts*. Princeton, NJ: Princeton University Press, 1992.
———. “Shared Intention and Personal Intentions.” *Philosophical Studies* 144, no. 1 (2009): 167–87.
Heersmink, Richard. “The Narrative Self, Distributed Memory, and Evocative Objects.” *Philosophical Studies* 175, no. 8 (August 2018): 1829–49.
Hofer, Myron A. “Hidden Regulators in Attachment, Separation, and Loss.” *Monographs of the Society for Research in Child Development* 59, no. 2/3 (1994): 192–207.
Huebner, Bryce. “Do You See What We See? An Investigation of an Argument Against Collective Representation.” *Philosophical Psychology* 21, no. 1 (February 1, 2008): 91–112.
———. “Genuinely Collective Emotions.” *European Journal for Philosophy of Science* 1, no. 1 (January 1, 2011): 89–118.
———. *Macrocognition: A Theory of Distributed Minds and Collective Intentionality*. Oxford University Press, 2013.
Hutchins, Edwin. *Cognition in the Wild*. Cambridge, MA: MIT Press, 1995.
Knorr Cetina, Karin. *Epistemic Cultures: How the Sciences Make Knowledge*. Cambridge, MA: Harvard University Press, 1999.
Krueger, Joel. “Extended Cognition and the Space of Social Interaction.” *Consciousness and Cognition* 20, no. 3 (September 1, 2011): 643–57.
———. “Ontogenesis of the Socially Extended Mind.” *Cognitive Systems Research*, Socially Extended Cognition, 25–26 (December 1, 2013): 40–46.
———. “The Who and How of Experience.” In *Self, No Self? Perspectives from Analytical, Phenomenological, and Indian Traditions*, edited by Mark Siderits, Evan Thompson, and Dan Zahavi, 27–55. New York: Oxford University Press, 2011.
LaBerg, David, and S. Jay Samuels. “Toward a Theory of Automatic Information Processing in Reading.” *Cognitive Psychology* 6 (1974): 293–323.
List, Christian. “What Is It Like to Be a Group Agent?” *Noûs* 52, no. 2 (June 2018): 295–319.
List, Christian, and Philip Pettit. “Aggregating Sets of Judgments: An Impossibility Result.” *Economics and Philosophy* 18, no. 1 (April 2002): 89–110.
———. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press, 2013.
Ludwig, Kirk. “Is Distributed Cognition Group Level Cognition?” *Journal of Social Ontology* 1, no. 2 (January 1, 2015).
Lyre, Holger. “Socially Extended Cognition and Shared Intentionality.” *Frontiers in Psychology* 9 (May 28, 2018): 1–9.
Mackenzie, Catriona. “Embodied Agents, Narrative Selves.” *Philosophical Explorations* 17, no. 2 (June 2014): 154–71.
Mackie, Diane M., Lisa A. Silver, and Eliot R. Smith. “Intergroup Emotions: Emotion as an Intergroup Phenomenon.” In *The Social Life of Emotions*, edited by Larissa Z. Tiedens and Colin Wayne Leach. Cambridge: Cambridge University Press, 2004.
Metzinger, Thomas. “The Myth of Cognitive Agency: Subpersonal Thinking as a Cyclically Recurring Loss of Mental Autonomy.” *Frontiers in Psychology* 4 (December 19, 2013).

———. “Why Are Dreams Interesting for Philosophers? The Example of Minimal Phenomenal Selfhood, Plus an Agenda for Future Research.” *Frontiers in Psychology* 4 (2013).
Minsky, Marvin. *The Society of Mind*. New York: Simon & Schuster, 1988.
Nagel, Thomas. “What Is It Like to Be a Bat?” *The Philosophical Review* 83, no. 4 (October 1974): 435–50.
Noë, Alva. *Out of Our Heads: Why You Are Not Your Brain, and Other Lessons from the Biology of Consciousness*. New York: Hill and Wang, 2009.
Overgaard, Søren, and Alessandro Salice. “Consciousness, Belief, and the Group Mind Hypothesis.” *Synthese*, March 1, 2019.
Pettit, Philip. “Groups With Minds of Their Own.” In *Socializing Metaphysics: The Nature of Social Reality*, edited by Frederick Schmitt, 167–93. New York: Rowman & Littlefield, 2003.
———. *The Common Mind: An Essay on Psychology, Society, and Politics*. New York: Oxford University Press, 1996.
———. “The Reality of Group Agents.” In *Philosophy of the Social Sciences*, edited by C. Mantzavinos, 67–91. Cambridge: Cambridge University Press, 2009.
Putnam, Hilary. “Psychological Predicates.” In *Art, Mind, and Religion*, edited by W. H. Capitan and D. D. Merrill. Liverpool: University of Pittsburgh Press, 1965.
Ridder, Jeroen de. “Representations and Robustly Collective Attitudes.” In *Socially Extended Epistemology*, edited by J. Adam Carter, Andy Clark, Jesper Kallestrup, S. Orestis Palermos, and Duncan Pritchard. Oxford University Press, 2018.
Roelofs, Luke. “Rational Agency without Self-Knowledge: Could ‘We’ Replace ‘I’?” *Dialectica* 71, no. 1 (2017): 3–33.
Rovane, Carol. *The Bounds of Agency: An Essay in Revisionary Metaphysics*. Princeton, NJ: Princeton University Press, 1998.
Rupert, Robert. “Minding One’s Cognitive Systems: When Does a Group of Minds Constitute a Single Cognitive Unit?” *Episteme* 1, no. 3 (February 2005): 177–88.
Salmela, Mikko, and Michiru Nagatsu. “Collective Emotions and Joint Action.” *Journal of Social Ontology* 2, no. 1 (2016): 33–57.
Schmid, Hans Bernhard. *Plural Action: Essays in Philosophy and Social Science*. Dordrecht: Springer, 2009.
———. “Plural Self-Awareness.” *Phenomenology and the Cognitive Sciences* 13, no. 1 (March 2014): 7–24.
———. “The Feeling of Being a Group.” In *Collective Emotions*, edited by Christian von Scheve and Mikko Salmela, 3–16. New York: Oxford University Press, 2014.
Schmitt, Frederick. “Group Belief and Acceptance.” In *From Individual to Collective Intentionality*, 61–96. Oxford: Oxford University Press, 2014.
Schwitzgebel, Eric. “If Materialism Is True, the United States Is Probably Conscious.” *Philosophical Studies* 172, no. 7 (2015): 1697–1721.
Searle, John. *Consciousness and Language*. Cambridge: Cambridge University Press, 2002.
———. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge and New York: Cambridge University Press, 1983.
Searle, John R. “The Intentionality of Intention and Action.” *Cognitive Science* 4, no. 1 (1980): 47–70.
Sender, Ron, Shai Fuchs, and Ron Milo. “Revised Estimates for the Number of Human and Bacteria Cells in the Body.” *PLoS Biology* 14, no. 8 (August 19, 2016).
Smith, Eliot R., Charles R. Seger, and Diane M. Mackie. “Can Emotions Be Truly Group Level? Evidence Regarding Four Conceptual Criteria.” *Journal of Personality and Social Psychology* 93, no. 3 (September 2007): 431–46.

Szanto, Thomas. “How to Share a Mind: Reconsidering the Group Mind Thesis.” *Phenomenology and the Cognitive Sciences* 13, no. 1 (March 2014): 99–120.
Thagard, Paul. *The Brain and the Meaning of Life*. Princeton, NJ: Princeton University Press, 2010.
Thalos, Mariam. *Without Hierarchy: The Scale Freedom of the Universe*. Oxford: Oxford University Press, 2013.
Theiner, Georg, Colin Allen, and Robert L. Goldstone. “Recognizing Group Cognition.” *Cognitive Systems Research,* Special Issue on Extended Mind, 11, no. 4 (December 1, 2010): 378–95.
Tollefsen, Deborah. *Groups as Agents*. Cambridge, MA: Polity Press, 2015.
———. “Organizations as True Believers.” *Journal of Social Philosophy* 33, no. 3 (Fall 2002): 395–410.
Tollefsen, Deborah, and Shaun Gallagher. “We-Narratives and the Stability and Depth of Shared Agency.” *Philosophy of the Social Sciences* 47, no. 2 (March 2017): 95–110.
Tononi, Giulio. “The Integrated Information Theory of Consciousness: An Updated Account.” *Archives Italiennes de Biologie* 150 (2012): 290–326.
Tononi, Giulio, and Christof Koch. “Consciousness: Here, There and Everywhere?” *Philosophical Transactions: Biological Sciences* 370, no. 1668 (2015): 1–18.
Verschoor, Stephan, and Bernhard Hommel. “Self-By-Doing: The Role of Action for Self-Acquisition.” *Social Cognition*, December 23, 2017.
Vold, Karina. “The Parity Argument for Extended Consciousness.” *Journal of Consciousness Studies* 22, no. 3–4 (2015): 16–33.
Wegner, Daniel M., Toni Giuliano, and Paula T. Hertel. “Cognitive Interdependence in Close Relationships.” In *Compatible and Incompatible Relationships*, edited by William Ickes, 253–76. New York: Springer Publishing, 1985.
Winnicott, Donald W. *The Child, the Family, and the Outside World*. Cambridge, MA: Perseus Publishing, 1987.