

Minding One's Cognitive Systems: When Does a Group of Minds Constitute a Single Cognitive Unit?

Robert D. Rupert, Texas Tech University

The possibility of group minds or group mental states has been considered by a number of authors addressing issues in social epistemology and related areas (Goldman 2004, Pettit 2003, Gilbert 2004, Hutchins 1995). An appeal to group minds might, in the end, do indispensable explanatory work in the social or cognitive sciences. I am skeptical, though, and this essay lays out some of the reasons for my skepticism. The concerns raised herein constitute challenges to the advocates of group minds (or group mental states), challenges that might be overcome as theoretical and empirical work proceeds. Nevertheless, these hurdles are, I think, genuine and substantive, so much so that my tentative conclusion will not be optimistic. If a group mind is supposed to be a single mental system having two or more minds as proper parts,¹ the prospects for group minds seem dim—or so I will argue.

I. Group minds, group states, and methodological guidelines

Three methodological principles guide the discussion to follow. First, any view according to which there are genuine group mental states (or cognitive states)² entails that there are group minds (or group cognitive systems), and thus cannot be evaluated independently of the hypothesis of group minds. One cannot simply assert the existence of genuinely autonomous group mental states—because, for example, that is what our everyday talk presupposes—while claiming that they are not states of minds. To do so would be to introduce a new notion of mental states, one deeply at odds with our normal understanding of them; that a *mental* state is the state of a *mind* is about as close as one gets to a definitional truth, this side of examples involving mathematics or kin-relations. (If there is any argument that convinces introductory philosophy students, it is Descartes's *cogito*: if there is thinking going on, a mind must exist.)

Secondly, we should address directly the issue of group minds. The nature of the mental (and the cognitive) has been extensively considered, independently of issues regarding group knowledge or collective action. We should not evaluate the hypothesis of group minds without taking these results into account. One might still wonder how we should pursue this evaluative tack. Margaret Gilbert worries that a critic might appeal to an allegedly essential property of minds to exclude group systems unjustifiably from the class of minds (Gilbert 2004, p. 10). This concern is duly noted: we should not assume too quickly that we know the *essential* properties of minds. All the same, there is, lurking in this vicinity, the possibility of *modus tollens*. A robust theory of group states entails that there are group minds (or, at least, group cognitive systems). If the entities to which we normally attribute group states (e.g., courts, corporations, unions) fail, by a significant number of established diagnostic measures, to count as minds, then this constitutes an exceedingly strong argument against the hypothesis of group mental states. Philosophers of mind and cognitive scientists have homed in on some of the central features of the mental and the cognitive and have provided promising, if not settled or fully satisfactory, accounts of some of these features: the capacity to represent, the flexibility of intelligence, the phenomenal quality of conscious experience. If we wish to know when a collection of two or more minds should be treated as a single mind, we should examine these composite systems to see whether they instantiate the central features of minds as we know them best. If group systems fail to exhibit a preponderance of these central features, then we have good reason to believe that these systems are not group minds. This form of argument is not demonstrative, but demonstration is not necessary to decisive refutation.

Third, I assume that the only effective way to support the hypothesis of group minds is to find distinctive causal-explanatory work for such entities to do. If our everyday ways of talking

seem to make reference to group states, yet our best accounts of human psychology and group systems do not appeal to group minds, I see no reason to accept the hypothesis of group minds. Given the naturalistic orientation of this essay, the present methodological guideline will play a significant role in what follows. Section II focuses specifically on general concerns that arise in connection with this third guideline, and section III invokes the guideline at various turns in the discussion of mental representation.

II. Groups, individuals, and causal-explanatory work

The hypothesis of *individual* minds clearly does causal-explanatory work. Human cognition exhibits high degrees of complexity and flexibility: humans reason, plan, and acquire skills in ways not determined by associative functions from past and current stimuli. Vague though this may be, it provides some indication of what a psychological theory must explain and why the hypothesis of an integrated system of internal representational states—a representational mind—seems so well suited to the task. Such a mind allows for complex, stepwise, internal processing, beyond the capacities of black boxes obeying only laws of association between stimuli and responses. Thus, cognitivism replaces behaviorism.

There will be more presently about causal-explanatory work, but first consider another common reason for rejecting a purely behaviorist conception of the individual: first-person experience. Individuals have conscious states, experienced from the first-person perspective. Each individual seems to have direct evidence of the existence of a mind distinct from behavior. There is no such evidence to support of the hypothesis of group minds. In the debate about group minds, there is nothing that plays the role occupied by individuals' claims to conscious experience in the debate over behaviorism. No group mind claims to have conscious experience. Hilary Putnam once suggested the possibility of super-stoics who could suppress all pain

behavior yet were keenly aware of their pain (Putnam 1975, p. 332). As a counterexample to behaviorism, the case seemed convincing enough to many of Putnam's readers, presumably because they could imagine themselves having the conscious experience of pain while imagining also their inability to move; but there are no analogous first-person reactions in the case of group minds.

The theory of consciousness is tricky business, one that I shall, for the most part, avoid. There are a couple of points to keep in mind, though. First, it seems that any form of consciousness that plays a causal role can be reduced to other cognitive or mental properties (e.g., properties having to do with access to information) (cf. Chalmers 1996, chapters 3 and 5). Given, then, the present emphasis on causal-explanatory work, it will be best to leave consciousness out of the discussion.³ Secondly, note that many philosophers and cognitive scientists have set aside questions about the phenomenal, irreducible aspects of consciousness when studying other aspects of the mind; and although one might rightly question whether the effects of this methodology have always been salutary, the strategy has by and large been fruitful. There remains, then, an interesting question whether group systems instantiate significant clusters of these other properties—the representational, computational, rational, perceptual, and architectural properties of minds. Given the central place consciousness holds in our conception of the mind, it will be better to talk only of group *cognitive* systems (and cognitive states) in the remainder. This will allow the discussion to go forward in a way that is charitable to the defenders of group states, while acknowledging that the apparent lack of group consciousness constitutes an initial strike against the hypotheses of group minds and group mental states.⁴

Return now to our concern with causal-explanatory work. What might be the systems relative to which group cognitive states do causal-explanatory work? Which groups or social institutions might be group cognitive systems? Common examples of such systems are labor unions, courts, and corporations. Although there are many details to be filled in, the advocate of group states might claim with some plausibility that such groups are cohesive representation-wielding systems instantiating something like intelligence. Less formal groups—a pair of people working together, a poetry reading group—are sometimes offered as candidates, but the more formal the institution, the more likely it is to possess the requisite degree of systemic coherence; thus, I will focus on the more formal institutions.⁵

One should also want to know what form such groups' cognitive states would take. The candidates most frequently suggested are expressions in a public language (Pettit 2003, p. 183, Gilbert 2004, *passim*; cf. Velleman 1997, p. 38, focusing on impromptu groups). Formal institutions typically have in place procedures governing the production of such representations. The canonical representations so produced might consist, for instance, in press releases or courts' written opinions. It seems explanatorily unnecessary to equate these physical formulations with autonomous cognitive states. After all, every step in the construction of such representations, as well as every step in the causal sequence alleged to involve the effects of those representations, proceeds either by brute physical causation (e.g., photons emitted from the surface of the page stimulate the reader's retinal cells) or by causal processes involving the mental states of individuals.⁶

Concerned to counteract a reductive approach to group states, Gilbert argues that it is neither a necessary nor sufficient condition for a group's being in cognitive state *C* that all or even a majority of its individual members be in *C* (Gilbert 2004, pp. 2-4). The existence of a

legitimate court decision, for instance, does not require that the opinion of the court match that of a majority of its members taken individually; nor does the fact that each member of a majority personally holds a given opinion entail that the court holds that opinion.

Gilbert's arguments are sound, I believe, yet they rebut only the simplest reductive views. It is misguided to think that a court's opinion is nothing more than a majoritarian agglomeration of the opinions of the court's members, set down on paper. Nevertheless, the construction of the court's opinion, its legitimacy, and its effect on society can all be explained without invoking anything beyond the conservative ontology of individuals and their states (together with photons, sheets of paper, ink, etc.). Bear in mind a few important facts about court opinions and the operations of courts (and note that similar points apply to other formally organized group systems). For a court's opinions to be issued in an orderly fashion that allows the continued existence and effectiveness of the institution, the justices, as well as members of the society in question, must have beliefs about what constitutes a canonical representation of the decision of the court. Furthermore, members of both groups typically have beliefs about the court's role in society as well as beliefs about the relationship between the court's issued opinions and the opinions of the court's members. As a final step in its examination of most of its cases, the U.S. Supreme Court issues a *majority* opinion, so labeled for a reason: it is assumed that the canonical representation of the court's ultimate opinion reflects the views held individually by a majority of justices. Odd cases where a majority of justices personally disagree with the majority opinion must be explained. A justice's willingness to vote against her settled individual view can be explained by her overarching theoretical motivations or her view about what it would be best for the court as a social institution to do given the particular political circumstances; however, a court in which the majority regularly disagrees with the court's canonical opinion, and where this

is widely known and is not accompanied by mitigating individual-level explanations, is a court not long of its place in society—and this lack of longevity would itself be best explained by the mental states of individuals, either members of society at large or members of whatever social institution is responsible for empowering the court.⁷ Thus, individual mental states explain why there are some departures from a simple majority-based aggregation, the way in which the process allows for these departures, and the reasons why these departures are not common. Furthermore, although canonical representations issued by Gilbert's court may have causal efficacy, they have it in virtue of individuals' states (together with nonpsychological processes). Posting group cognitive states of the court does no causal-explanatory work, for there is none left to do.⁸

My passing mention of brute-physical processes might strike a critical note in the minds of some readers, raising the concern that a double-standard drives the preceding deflationary argument. The argument freely invokes individual mental states, yet individuals' mental states are identical to, realized by, or supervene on physical processes. If we are to dismiss group cognitive systems on the grounds that they have no distinctive causal-explanatory work to do, why not dismiss individual minds as well? Since we emphatically should not omit individual minds from our ontology, our objector will conclude, neither should we omit group cognitive systems.

This objection has little merit, for the two cases differ greatly with respect to our understanding of what are sometimes called 'inter-level relations'. We only vaguely understand the connection between the physical basis of individual mental states and the states themselves. This is especially pointed with regard to the psychological regularities formulated at the level of the individual: we have no idea how to reduce these psychological regularities (or laws), even

schematically, to physical laws involving, say, brain states; neither are we in any position to eliminate, even in principle, talk of individual mental states in favor of talk about brain states (although I do not mean to imply that neuroscience has been unproductive).⁹ As the preceding discussion of courts should illustrate, though, we understand well the relationship between canonical representations formulated by groups or their appointed agents and the mental states of the individuals involved. Granted, the relation is more complex than that of an aggregate to its parts; nevertheless, there is no reason to exclude consideration of the more complex, but well enough understood, network of individual states that explain the issuing of and effectiveness of group canonical representations.

III. Representation in group systems

It is a central feature of cognitive systems that they possess representational capacities: cognitive states typically correspond to, refer to, or are about something in the world beyond the subject's mind. This aspect of cognition has been widely discussed and a number of promising theories of mental representation have emerged.¹⁰ In what follows, I survey five naturalistic approaches to mental representation. None of the theories applies in a natural or convincing way to group states. This result threatens the hypothesis of group cognitive states, not only because representation and cognition are so tightly connected in the cognitive sciences, but because going theories of group cognitive states themselves emphasize the role of representation in group cognition: group states appear in concrete form as, for example, press releases and written opinions.

Although the states in question take what would appear to be linguistic form, they are supposed to be *cognitive* states. Thus, the representations that encode these states should not be approached as bits of public language; they are supposed to be literally thoughts, not words. I

assume this view throughout the remainder of this section, for it seems to be the only view that takes seriously the hypothesis of group cognitive states. Doing so creates a problem, however, by unnecessarily complicating the theory of meaning. If the hypothesis of group cognitive states is correct, then a theory of *mental*, not linguistic, content should apply to group deliverances. This conflicts, though, with the obvious explanation of the meaning of, say, a court's written decision. Being a piece of public language, one would expect a theory of linguistic meaning (almost certain to be distinct from a theory of mental content) to account for its content. It is quite implausible to say that a bit of public language, an English sentence, for example, has one meaning in a group's press release, but a different meaning when uttered by an individual— simply because in one case a group is involved and in the other case, an individual speaker (this is especially troubling when the same person is speaking in both kinds of case, sometimes as press agent, sometimes as individual); yet, this seems to be just what the advocate of group states must say.¹¹

Let us set aside this general concern and consider the application of particular theories of mental content to group states:

Indicator semantics: Fred Dretske (1988) proposes that a mental representation *R* represents that property *P* is instantiated if and only if *R* has an acquired function within the cognitive system (control over a particular kind of movement, for example) and it acquired that function because it indicated *P*. One thing's indicating another can be spelled out in terms of conditional probabilities: *A* indicates *B* if and only if the probability of *B* given the occurrence of *A* is one (and *B*'s occurrence is not logically necessary).

Problems for the application of indicator semantics to group systems:

1. Group systems appear to lack indicating structures. In the individual subject, there are, early in development, perceptual structures that indicate (or do something less demanding than indication—see Slater 1994, Maloney 1994), and it is plausible that these acquire control over bodily movements in a way that is best explained by the indicational properties of those perceptual structures. In the standard group system, a court or a corporation, for instance, there do not appear to be analogous structures, i.e., structures that indicate the presence of various kinds or states in the environment such that they acquire new roles in the group system as a result of their indicational powers.

2. Any semantics of mental representations must allow for the flexible use of representations; it must leave room for mental representations to be used in contexts other than perception or detection. The indicational framework can best accommodate this demand by allowing mental representations to become detached from perceptual processing, while inheriting their content from the content of the perceptual states that cause them (or are related to them in some other privileged way). Here we have something like a distinction between representations that appear in a perceptual module and representations used in central processing (Fodor 1983). There is little reason to think that a standard group's architecture includes this division. My concern here is not that a group system must have precisely the same architecture as our paradigm minds, only that group systems possess the sort of architecture that allows our best theories of representation to apply to them. This seems to require a distinction between indicational structures to which representational content is initially fixed and content-bearing structures that are not part of a perceptual or quasi-perceptual systems. (Although I will not repeatedly raise this concern about perception and group cognitive architecture, it is important to note that this concern could be raised with respect to all of the theories of representation discussed below, most clearly to the

pure-informational and causal-historical theories, but to the others as well on at least one of their interpretations considered—this includes the conceptual-role theory addressed only in note 17.)

Pure-informational semantics: Jerry Fodor (1987, 1990) is the foremost proponent of this view.

On one formulation, R represents P if “ P s cause R s” is a law and any other law-like relation between R and its causes is asymmetrically dependent on P s’ causal relation to R s; for some Q s’ causal relation to R s to be asymmetrically dependent on P s’ causal relation to R s is for it to be the case that in the nearest possible worlds (roughly, the merely possible situations most similar the actual world) where Q s’ nomic relation to R s is severed, P s cause R s all the same, but not vice versa. The basic idea here is that, although things other than P can cause R , their causing of R is dependent on the relation between P and R , but the relation between P and R does not depend on the relation between R and any of those other things.

Problems:

1. It is difficult to see why the proper causal relations would hold independently (or independently enough) of the asymmetric dependencies into which individuals’ mental representations enter. It is plausible that ‘speech’, as that word might appear in a court decision regarding freedom of speech, is such that all of its causes besides actual speech are, in their ability to cause ‘speech’, asymmetrically dependent on the causing of ‘speech’ by speech. This is plausible, however, because it follows straightway from the fact that certain individuals have mental representations with the relevant contents. We would like a theory of mental content that attributes content, in the first instance, to representations that are part of a group cognitive system. Otherwise, there is a good case to be made that group cognitive systems, *qua* group systems, contain no representations whose content is not derived from the content of

representations in some other system—not what we should want in a genuine cognitive system (or mind).¹²

Teleological semantics (teleosemantics):

The fundamental idea behind teleosemantic theories is this: a structure R represents an individual, kind, or property P if and only if R 's bearing some privileged relation to P accounts for the continued reproduction of R . Some such views emphasize the evolutionary selection of R s (Millikan 1984), while others focus on the current maintenance of R (Schlosser 1998).

Problems:

1. Conceived of as strictly evolutionary theories of mental content, teleosemantic views simply cannot apply to the groups in question. There is no straightforward way to apply evolutionary theory to such entities as courts or corporations. Such groups and their ancestors have not existed for long enough; what is more, given the lack of group DNA there would appear to be (1) nothing that varies in such a way that its differences might then be inherited by the group and (2) nothing to encode successful variations in fitness, so that they may be passed on to descendants. Nevertheless, teleosemantics admits of a nonevolutionary reading. How do the representations in group systems fare when teleosemantics is interpreted in this way?
2. Here we confront a complication within teleosemantics itself, which, once taken into consideration, leads to a problem much like one we have already encountered. Nonevolutionary versions of teleosemantics focus on the relation a symbol currently bears (or recently bore) to something in the environment, which relation accounts for the continued production or maintenance of the symbol (say, a neural structure) of interest. Typically this involves perceptual recognition or some other form of causal relation (Millikan identifies this relation with that of carrying information about—1984, p. 146). Trouble now arises for the advocate of

group states, because there would appear to be no plausible account of the relevant causal mechanism that does not simply run through and exploit the causal correspondences between individuals' mental representations and what they represent.

Causal-historical semantics:

As this view has it, mental content is, at root, causal-developmental (Rupert 1998, 1999, 2001, Prinz 2002). It is likely that evolutionary processes selected for neural mechanisms that allow individuals' mental representations to enter into useful causal relations to properties, kinds, and individuals in the environment. Nevertheless, to represent is to enter into the relevant causal relations. Like most neural structures, the cohesive neural units that constitute (or realize) mental representations are shaped developmentally, by interaction with the environment. On the view I advocate, this shaping involves a certain statistical pattern of interaction with the very things that thereby come to be represented. Prinz is more inclined toward a "one-shot" approach (Prinz 2002, chapter 9), according to which the initial cause of a mental representation then constitutes its content (other conditions must also be met, on Prinz's view, but these are much like Fodor's pure-informational conditions, already discussed).

Problems:

The lack of distinctive perceptual capacities precludes application of a causal-historical theory to group cognitive systems. On my view, causal interaction with the environment shapes the neural resources of the individual subject, resulting in mental representations that bear a privileged relation to what they represent. Prinz describes his view as an *empiricist* theory of concept formation, and thus one that depends very heavily on sensory interaction with the environment. Where in a group system's interaction with the environment can we locate such processes?

Teleo-isomorphic semantics:

What I will call the ‘teleo-isomorphic’ view grounds mental representation in a match between the structural properties of representations and the structural properties of what they represent: R represents P because the structure of R mirrors the structure of P . Robert Cummins (1996) has proposed the most prominent version of this account, and the version on which I will focus.

According to his view, isomorphic relations among structures determine the contents of mental representations. Isomorphic relations are too common, however, to fix fully the contents of our thoughts: a single representing structure can be isomorphic to a number of other structures that we would not take to be the objects of a thought involving that mental representation (Cummins 1996, p. 120). To avoid rampant and unacceptable ambiguity of thought content, this approach must incorporate some further, winnowing principle. Cummins adverts to the *function* of the portion of the cognitive system in which R appears. For example, a collection of co-firing neurons in visual cortex will almost certainly be isomorphic to many things; given, though, that the function of (at least certain parts of) visual cortex is to represent facts about the subject’s immediate environment, the teleo-isomorphic view can avoid the embarrassment of having to assign all of those many things as the contents of the subject’s perceptual state.¹³

Problems:

1. Teleo-isomorphic semantics seems best suited to cases of perception and motor control. It seems clear enough how a set of neurons can structurally mirror aspects of the immediate physical environment and how their doing so would facilitate the subject’s navigation of that environment. (Cf. Cummins’s example of an autobot—a small, automated car—that runs a maze using a card the slots in which are isomorphic to the maze’s clear path; 1996, pp. 94ff.) For obvious reasons, however, the advocate of group mental states should not want to view the teleo-

isomorphic theory in this way; in the typical group case, there simply is no subject navigating the environment.¹⁴

2. Cummins presents his theory as a picture theory of representation, which would seem to limit mental representation to the representation of objects and scenes in concrete physical space (of the sort one could convey realistically in a painting or photograph); such an interpretation would be misleading, though, given that the picturing relation Cummins has in mind—*isomorphism*—is an abstract logical relation, applying to mathematical and theoretical contexts as much as it does to spatial and perceptual ones. All the same, this increased generality does not seem to help the advocate of group states. Representations alleged to constitute group states are typically given linguistic expression or expression in some other conventional symbol system. The ink marks that constitute a court's decision are, taken as an entire structure, isomorphic to some other structures, but we have no reason to think that the abstract structures the decision is about—abstract conceptual structures, typically—will be among those things structurally mirrored by the arrangement of ink on page. This is precisely where Cummins takes his own approach to break down (Cummins 1996, p. 111). As a result, Cummins offers a distinct account of meaning for public language.¹⁵

3. Recall the reason for introducing the teleological aspect of the view under consideration. Of the many things a given *R* is isomorphic to, a particular thought (belief, intention) involving *R* is typically about only one of these things because there is some sense in which it is the function of the part of the cognitive system in which *R* appears to represent things of a certain kind. The appeal to teleology introduces a further difficulty for the hypothesis of group states; for, although it seems plausible enough that certain products of group processes (e.g., the pronouncements of certain group members in their roles as administrators or press agents) have the function of being

about certain kinds of thing rather than other kinds, it does seem that our best account of why those parts have those functions adverts in a straightforward way to the mental states of individuals. Group members intend that their agents talk about certain matters rather than others, either in general—because certain agents have particular ongoing responsibilities—on an *ad hoc* basis.¹⁶

In summary, advocates of group states face significant hurdles with respect to the theory of representation. What are alleged to be group cognitive states appear to take concrete form as structures in a public language. When viewed this way, the semantics of such structures is almost certainly different from the semantics of cognitive structures. When we attempt to treat these structures instead as genuinely cognitive and apply theories of mental representation to them, numerous difficulties arise. In many of the cases surveyed, a theory of mental representation seems to apply to group states only in virtue of the contents of individuals' mental representations. Furthermore, our best accounts of mental content seem to exploit architectural features that groups typically lack: many of our best explanations of how mental representations get their content assign a privileged role to perceptual or quasi-perceptual processing, thereby requiring a cognitive architecture that group systems typically do not possess. Problems specific to the application of the various theories of content further dim the hopes that group systems, considered as genuinely autonomous cognitive systems, might comprise representational states.¹⁷

IV. The frame problem and the global nature of cognition

Let me close by identifying one further point of disanalogy between cognition as it is best known and the features of typical group systems. The global nature of certain aspects of cognition has made quite an impression on many cognitive scientists, most notably via the frame problem in A.I. Put roughly, any belief (or similar cognitive state such as a memory) can be relevant to the

evaluation of any of the subject's other beliefs. Furthermore, properties of large sets of the subject's beliefs or ways in which members of those sets are interrelated often determine the relevance of a given cognitive state to another (Fodor 2000, chapter 2). There might in fact be many frame problems and many morals; clarifying these matters is beyond the scope of the present discussion. It is, however, provocative to note the difference between the global nature of cognition in the individual and the not-so-global nature of what is alleged to be cognition in group systems. In typical group systems, there are clearly defined channels of access and manipulation of the representations that are supposed to be groups' cognitive states. To the extent that there is a global aspect to group processing, it is introduced by the individual participants. Group organization does nothing to facilitate global cognitive processing, except insofar as group organization makes a place for individuals who can engage alone in global cognition and inject their results into the group process. There is no independent frame problem for groups, and this should give further pause to advocates of group states.

Works Cited

- Adams, F., and K. Aizawa (2001). "The Bounds of Cognition." *Philosophical Psychology* 14: 43-64.
- Block, Ned. (1986) "Advertisement for a Semantics for Psychology." In P. French, T. Uehling, and H. Wettstein (eds.), *Midwest Studies in Philosophy*, Vol. 10: *Studies in the Philosophy of Mind* (Minneapolis: University of Minnesota Press), pp. 615-78.
- Chalmers, D. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.
- Chalmers, D. (2002). "Consciousness and Its Place in Nature." In D. Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings* (Oxford: Oxford University Press), pp. 247-72.
- Cummins, R. (1996). *Representations, Targets, and Attitudes*. London: MIT Press.
- Dretske, F. (1988). *Explaining Behavior: Reasons in a World of Causes*. London: MIT Press.
- Fodor, J. (1983). *The Modularity of Mind*. London: MIT Press.
- Fodor, J. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. London: MIT Press.
- Fodor, J. (1990). *A Theory of Content and Other Essays*. London: MIT Press.
- Fodor, J. (1994). *The Elm and the Expert: Mentalese and Its Semantics*. London: MIT Press.
- Fodor, J. (2000). *The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology*. London: MIT Press.
- Fodor, J., and E. LePore (1992). *Holism: A Shopper's Guide*. Oxford: Blackwell.
- Gilbert, M. (2004). "Collective Epistemology." *Episteme* 1: 0-17 (pagination refers to typescript).

- Goldman A. (2004). "Group Knowledge Versus Group Rationality: Two Approaches to Social Epistemology," *Episteme* 1: 11-22.
- Grush, R. (1997) "The Architecture of Representation." *Philosophical Psychology* 10: 5-23.
- Hutchins, E. (1995). *Cognition in the Wild*. London: MIT Press.
- Maloney, J. C. (1994). "Content: Covariation, Control, and Contingency." *Synthese* 100: 241-90.
- Millikan, R. (1984) *Language, Thought, and Other Biological Categories: New Foundations for Realism*. London: MIT Press.
- Pettit, P. (2003) "Groups with Minds of Their Own," in F. Schmitt (ed.), *Socializing Metaphysics*, Oxford: Rowman and Littlefield.
- Prinz, J. (2002). *Furnishing the Mind: Concepts and Their Perceptual Basis*. London: MIT Press.
- Putnam, H. (1975). *Mind, Language, and Reality: Philosophical Papers, Volume 2*. Cambridge: Cambridge University Press.
- Rupert, R. (1998). "On the Relationship between Naturalistic Semantics and Individuation Criteria for Terms in a Language of Thought." *Synthese* 117: 95-131.
- Rupert, R. (1999). "The Best Test Theory of Extension: First Principle(s)." *Mind and Language* 14 (1999): 321-355.
- Rupert, R. (2001). "Coining Terms in the Language of Thought: Innateness, Emergence, and the Lot of Cummins's Argument against the Causal Theory of Mental Content." *The Journal of Philosophy* 98: 499-530.
- Schlosser, G. (1998) "Self-reproduction and Functionality: A Systems-Theoretical Approach to Teleological Explanation." *Synthese* 116: 303-354.
- Segal, G. (2002) *A Slim Book about Narrow Content*. London: MIT Press.

Slater, C. (1994). "Discrimination without Indication: Why Dretske Can't Lean on Learning."

Mind and Language 9: 163-80.

Velleman, J. D. (1997) "How to Share an Intention," *Philosophy and Phenomenological*

Research 57: 29-50.

Wilson, R. (2004) *Boundaries of the Mind: The Individual in the Fragile Sciences*. Cambridge:

Cambridge University Press.

¹ Robert Wilson allows that there could be a group mind none of the proper parts of which is an individual with a mind (Wilson 2004, p. 282). Such a mind, though, would appear to be an individual mind, plain and simple, of just the sort that humans have: a mind made up of (or realized by, or supervening on, or what have you) a collection of parts, that do not themselves have minds. Wilson describes the possibility the way he does because he is specially interested in the case of insect colonies, where the group systems in question consist of biological, not necessarily psychological, individuals (Wilson 2004, 291-92); but perhaps here is a point where the interests of the social sciences and the biological sciences diverge more than Wilson's unified approach can comfortably acknowledge.

² The importance of including 'cognitive' as an alternative to 'mental' will presently become clearer.

³ I hesitate to oversimplify these matters too much, though. One can, I think, make a good case for categorical properties that serve as the ultimate grounds of causal relations, and do so even though neither the categorical properties nor their instantiations enter directly into causal relations. If, as Chalmers (1996, 2002) suggests, phenomenal consciousness is the instantiation of certain of these ultimate categorical bases, phenomenal consciousness will, in some looser but finally more germane sense, play a causal-explanatory role: its presence will provide the metaphysical ground, however precisely that is to be spelled out, of causally efficacious properties.

⁴ Wilson frames his discussion of group minds in similar terms and does so for similar reasons (Wilson 2004, p. 290). Present standards of ontological inference are somewhat more demanding than Wilson's, however. Wilson is willing to consider a thesis of minimal group minds, such that the defender of group minds will have made significant progress if a group appears to be the legitimate subject of even a single psychological property. This is, however, too slim a ground for ontological inference. More likely, if a system appears to instantiate only one psychological property—say, one belief-like state on one occasion—this appearance warrants only an as-if attribution of mentality or cognition (and this is true even if, more charitably, we take Wilson to mean one *kind* of psychological property, rather than a single instantiation of a psychological property).

Let me briefly bring out a further disanalogy between group systems and conscious minds. It is often thought that a mental representation of the self plays a special role in the life of a mind, particularly in self-consciousness. Admittedly, many group systems have names, written on letterhead or painted on signs. It is another matter, though, to show that such representations play the role of a concept of one's self.

⁵ Here are some more detailed reasons for focusing on formally organized groups. Our everyday ways of talking provide the primary motivation for attributing mental states to groups organized less formally. Given, though, that our everyday ways of speaking also cut strongly against the existence of group minds, it is more to the point to limit the discussion to cases where the attribution of group states does some causal-explanatory work. Potential for such work seems clearest in the case of the formally organized social institutions studied by social scientists. Of course, such scientists do sometimes study informally organized groups, but many of these—mobs, for instance—have no claim to the systemic unity of a cognitive system, and so should be set aside. Thus arises a second reason for limiting the discussion to formal institutions. Many kinds of thing have persistence conditions, i.e., a requirement that their instantiations exist for a minimum amount of time; consider the kinds *mountain*, *epidemic*, and *stain*. The ability to learn is a central feature of cognitive systems, and this requires that they persist through time and maintain a fairly high degree of structure throughout that time. Formally organized groups are the ones most likely to meet these requirements.

I do not mean to take too hard a line here. Whether a group has the requisite temporal and structural coherence depends on the details of the case, as is illustrated by Gilbert's example of the long-running poetry reading group to

which she belongs (Gilbert 2004, p. 3). I should note, however, that the more unity a group possesses, the more likely it would seem that the concerns expressed below about formally organized groups will apply to it. Group coherence has to be achieved somehow; the best, perhaps the only way, to effect such coherence is through the sorts of mechanisms that give rise to the concerns in question.

⁶ Cf. Wilson's discussion of the social manifestation thesis, which he offers as a contrasting, and less metaphysically robust, explanation of much of the data thought to support the hypothesis of group minds (Wilson 2004, chapters 11 and 12).

⁷ An interesting side note: There is currently a debate raging in the state of Texas over the Texas state legislature's policy of not recording the votes individual legislators cast on bills under their collective consideration (it seems unimaginable to any advocate of good government that such a record is not currently kept, but there you have it). Many residents of Texas are incensed by the situation (see, for example, the editorial and letters page of the *Dallas Morning News*, 12/27/2004), not being at all satisfied with a canonical representation of the will of the legislature as a whole; these citizens wish to know precisely which individuals voted which way regarding every piece of proposed legislation on which a vote is taken.

⁸ An analogous story is easily told about Philip Pettit's examples of collective or group rationality (Pettit 2003). Individuals do not want their institutions to become ineffective or distrusted. Thus, their individual mental states together explain why they choose the "premise-centered procedure" (Pettit 2003, p. 168) and why they employ *modus ponens* reasoning over time (Pettit 2003, p. 174) to produce the group's decisions.

⁹ There are further complications attached to the full-out elimination of cognitive states, some having to do with the apparently self-defeating nature of a scientific theory that denies all representation and cognition. I am not claiming that the eliminativist cannot handle such complications—that remains to be seen—only that elimination at the individual level faces hurdles beyond those faced at the group level; this provides an additional reason to think that my argument against autonomous group states cannot simply be carried over as an eliminativist argument at the individual level.

¹⁰ Although I continue to talk of 'cognitive' states and systems, I use 'mental' to talk about representations and their content. 'Cognitive content' is not commonly employed in discussions of mental representation, and when it is used, it is typically to label a different sort of content from the kind I am interested in here, something more like what philosophers have called 'narrow content' (see, e.g., Prinz 2002, chapter 10; Segal 2000, p. 4).

¹¹ This problem could be solved by adopting the view that group cognitive states are represented only linguistically. Then, however, the advocate of group states would face the difficult task of explaining why the meaning of such linguistic units is not simply derived from the content of the mental states (or linguistic usage) of individuals. (Cf. Adams and Aizawa's discussion of derived and nonderived meaning when criticizing the view that individual minds extend into the environment; 2001, pp. 48-49.)

¹² Fodor himself allows intentional mediation to ground asymmetric dependencies. This is not the same, however, as requiring content to be fixed in proper parts of a cognitive system (the proper parts being individual group members) before any content can be fixed for any representation in the cognitive system in question (in this case, the group system). (Complications arise regarding Fodor's proposed solution to Quine's problem of the indeterminacy of reference [Fodor 1994]. Nevertheless, Fodor's manner of resolving indeterminacy—by assuming that conceptual roles fix the meaning of logical connectives—offers nothing that the advocate of group states could utilize to solve the problem raised in the text.)

Also note that, although the pure-informational view might seem to be the most promising theory of content for group states, it has been the object of intense criticism; the fourth and fifth theories of content discussed herein were, to a significant extent, introduced as ways to get around these criticisms.

¹³ Rick Grush (1997) advocates a view similar in many ways to Cummins's. They disagree about the nature of the structural mirroring: in contrast to Cummins's appeal to isomorphism, Grush emphasizes counterfactual dynamical and functional equivalence between representing structures and what is represented (on his view, a representing structure is used to emulate the behavior of a real-world system that has the same dynamical or functional structure). Furthermore, according to Cummins, teleology plays the disambiguating role discussed in the main text, whilst Grush claims that the way the subject *uses* a representational structure does the necessary work (Grush 1997, pp. 17, 19). Nevertheless, in the present context, the common underpinning of Cummins's and Grush's views—that representation is grounded in structural mirroring—provides a fine enough view. (Note also that Cummins sometimes seems concerned specifically with dynamical relations, in particular with the exploitation in reasoning of dynamical change in the structures of the mental representations being used as proxies for what they represent; see Cummins 1996, pp. 95, 110-11.)

¹⁴ The navigational teams discussed by Hutchins (1995) might be an exception here; see note 15 for more discussion of this case.

¹⁵ The present objection does not apply to a group system that produces *maps* instead of sentences—some of what Hutchins’s navigational teams produce seems to fit this description. Still, for this fact to count as a genuine success for the advocate of group states, other objections to the teleo-isomorphic approach, as well as our more general objections, must be met. For example, even if the map produced by a group system is isomorphic to the correct geographical structure, this isomorphism effectively guides the pursuant trajectory of the ship not by a motor control system but because members of the crew individually believe that the picturing relation holds between their collectively constructed map and the environment, and they act accordingly. Furthermore, the map has the particular function of representing the immediate geography, rather than something else to which it is isomorphic, only because the relevant individuals take it to do so or mean for it to do so.

¹⁶ This talk of a part of the cognitive system’s having the function of representing something oversimplifies Cummins’s account by failing to separate intenders, representations, and attitudes in the rich way integral to Cummins’s full-blown theory. All the same, rectification will be of no service to the advocate of group states. The advocate of group states will have no easier time explaining why it is the function of intenders to take certain targets than she will have explaining why it is the function of a certain part of the system to represent certain kinds of thing.

¹⁷ Absent from this discussion has been any consideration of conceptual- or functional-role semantics. The most plausible theories of mental content that invoke conceptual- or functional-roles are two-factor theories (Block 1986), where the factor other than conceptual role is one of the factors discussed in the main text; thus, a plausible conceptual-role account of mental representation will inherit some of the problems already surveyed, if it is pressed into service by advocates of group states. In contrast, *pure* conceptual- and functional-role theories seem to me to be hopeless nonstarters as realist theories of mental representation. (Support for this pessimistic view can be found in Cummins 1996 and Fodor and LePore 1992.)