

This is the penultimate version of a paper forthcoming in *Phenomenology and the Cognitive Sciences*. The final publication is available at www.springerlink.com.

**Thinking Things and Feeling Things:
On an Alleged Discontinuity in Folk Metaphysics of Mind[†]**

Mark Phelan
Lawrence University

Adam Arico
The University of Arizona

Shaun Nichols
The University of Arizona

How do people ordinarily attribute mental states to other entities? Clearly, people take physical features into account when assessing whether an organism is likely to occupy particular mental states. An eyeless cave fish, for instance, will be thought unlikely to occupy visual states. However according to one recent theory, people use information about physical constitution not only in this piecemeal fashion to determine which mental states an organism is likely to occupy, but also to draw a fundamental distinction between entities that can merely think and entities that can also feel (Knobe & Prinz 2008). According to this view, people recognize a deep *discontinuity* between phenomenal and intentional states, such that they refrain from attributing feelings and experiences to entities that do not have the right kind of body, though they may attribute thoughts to entities that lack a biological body, like corporations, robots, and disembodied souls. Alternatively, some have denied that there is any deep discontinuity between the physical features that lead us to attribute the two varieties of mental states (Arico et al. 2011).

[†]Portions of this paper were presented at Brown University's Social Cognitive Science Research Center, the European Workshop on Experimental Philosophy at Eindhoven, Netherlands, the London School of Economics' Philosophy Department, the Metro Experimental Research Group, the Southern Society for Philosophy and Psychology, Yale University's Experimental Philosophy Lab, Yale University's Mind and Development Lab, and the University of Arizona's Experimental Philosophy Lab. Audience comments helped improve this paper. The authors are also grateful to Michael Bruno, Ben Chan, Georg Kjøll, Joshua Knobe, Eric Mandelbaum, Justin Sytsma, Zoltán Gendler Szabó, Jonathan Weinberg, and three blind referees for this journal who gave valuable comments on earlier drafts.

In other words, the cognitive process that leads us to attribute mental states to various entities does *not* distinguish between the physical features necessary for intentional states and the physical features necessary for phenomenal states.

In this paper we examine some of the research that has been used to motivate the discontinuity view. Specifically, we focus on experiments that examine people's intuitions regarding the aptness of various mental state ascriptions to groups. These studies have been taken as evidence that people are more inclined to think of groups as having intentionality than as having phenomenology. This result, combined with the fact that groups obviously lack a single biological body, has been taken as evidence that people use information about physical constitution in fundamentally different ways when attributing the two kinds of states. However, as we will explain, these studies support a discontinuous picture of folk metaphysics of mind only on the assumption that the experimental participants are interpreting the relevant group mental state ascriptions in a very specific way. Thus, we empirically investigate how people are interpreting group mental state ascriptions and present evidence that they are *not* interpreting these ascriptions in a way that supports the discontinuity view. Instead, we will argue that people generally interpret group mental state ascriptions distributively, as attributions of mental states to various group members.

Constructing theories of folk psychology based on how people talk about minds and mental states is a common method within philosophy and psychology. However, as our discussion of research into the folk psychology of group minds will illustrate, this method often rests on specific assumptions about how experimental participants are interpreting the sentences they speak and endorse. As we will argue, such assumptions are not always warranted.

1. The Discontinuity Thesis

Many philosophers (e.g., Block 1995, Nagel 1974, Rosenthal 1997) draw a sharp distinction between mental states that feel a certain way – *phenomenal* states – and those that don't. In particular, philosophers often maintain that phenomenal states like pain and pleasure can be distinguished from states of believing, desiring, intending, and other *intentional* states. Some philosophers and psychologists have recently argued that ordinary people also appreciate this key distinction between phenomenal and intentional states (e.g., Knobe and Prinz 2008; Gray, Gray, and Wegner 2007; Robbins and Jack 2006).¹

In this paper, we want to examine some evidence that seems to support the idea that the physical features that typically dispose people to attribute intentional states to an entity differ from the physical features that typically dispose people to attribute phenomenal states to an entity. Call this the *discontinuity thesis*.² It can be contrasted with what we will call the *continuity* view, according to which the sorts of physical properties that dispose people to attribute intentional states to an entity also dispose people to attribute phenomenal states to that entity.³

Now, it's rather plausible that there really is a discontinuity between what is required for the capacity to have intentional states and what is required for the capacity to have phenomenally conscious states. But the discontinuity thesis concerns our everyday, common-sense psychology, not metaphysical fact. And while learning about folk judgments won't settle the metaphysical

¹ Of course it's controversial whether all intentional states are non-phenomenal. The thesis under consideration must maintain that at least a core set of intentional states are.

² The discontinuity thesis implies a second claim about folk psychology. In claiming that cognition relies on different sets of physical features for the distinct kinds of mental state attributions, the discontinuity thesis implies the additional claim that ordinary people distinguish (perhaps tacitly) phenomenal from (non-phenomenal) intentional states. This second claim has been contested, both empirically and theoretically, elsewhere (Arico 2010, Sytsma and Machery 2009, 2010). However, in this paper we remain neutral on this issue.

³ One example of a continuity view can be found in Arico et al. (2011), which defends a cognitive model of mental state attribution dubbed 'the AGENCY Model'. On that model, our everyday, intuitive attributions of intentional states and phenomenal states are *both* consequences of categorizing a thing as an 'AGENT'.

question of whether there is a genuine discontinuity, it may nonetheless illuminate other philosophically relevant questions. For instance, a number of researchers have recently suggested that attributions of conscious states and attributions of intentional states play fundamentally different roles in moral cognition (e.g., Knobe and Prinz 2008; Gray, Gray, and Wegner 2007; Robbins and Jack 2006). Thus, it is possible that the question of whether the folk, in fact, see a discontinuity between what can occupy phenomenal states and what can occupy intentional states potentially has direct bearing on our understanding of the psychology behind our moral judgments and practices. The study of lay judgments concerning phenomenal states might also contribute to issues, both contemporary and historical, in philosophy of mind. For instance, a number of theorists have recently maintained that some intentional states are constitutively determined by their phenomenal character (Kriegel 2003, Horgan and Tienson 2002, Strawson 2005), and the discovery that people tend to reject discontinuity might shed light on debates surrounding this 'phenomenal intentionality' thesis. Likewise, the work on everyday intuitions about consciousness attributions might also help us assess the extent to which we should take those intuitions to inform our metaphysics of mind (see, e.g., Fiala et al. 2011).

One way that the discontinuity thesis might gain support is if theorists found that people are willing to grant that certain entities are capable of one kind of state but not the other. This is precisely the strategy taken by Gray, Gray, and Wegner (2007). They had experimental participants make pair-wise comparisons of different characters' capacities for different forms of mental activity. For example, they asked participants whether a chimpanzee or a man in a persistent vegetative state was more capable of feeling pain. After analyzing participants' responses to these comparisons, Gray et al. contend that people attribute mental states along two

dimensions: “Agency” and “Experience”.⁴ According to Gray et al., people regard entities as varying in their capacity for *Agential* states (such as thinking, planning, and communicating) and *Experiential* states (such as feeling hunger, fear, and pain). This seems to correspond roughly to the intentional/phenomenal distinction. Additionally, Gray and colleagues maintain that some kinds of entities are seen as having one kind of state but not the other. For instance, they report that people regard God as having high levels of Agency but low levels of Experience. Gray, Gray, and Wegner (2008) have also found that people treat Google in the same way as God, ascribing high levels of Agency but low levels of Experience. Fetuses and frogs, by contrast, are taken to have much Experience but little Agency.

Although this is evidence for some tendency to attribute certain mental states to one entity rather than another, it is not clear that it constitutes evidence for the discontinuity thesis that we’re exploring. For Gray and colleagues’ category of Agency includes mostly very complex capacities (e.g., self-control, morality, planning). These capacities are clearly much richer than the simple notion of an entity that has some intentional states (e.g., a desire). On the other hand, the category of Experience includes mostly experiential states (e.g., hunger, fear, and pain)—states that are inappropriate for God and Google to possess for reasons other than any lack of mental capacity. What, for instance, would God be afraid of? As a result, the evidence from Gray et al doesn’t necessarily show that people think certain entities have intentional states but not phenomenal states. Just because certain entities ranked low on *these* phenomenal capacities doesn’t mean they are regarded as lacking phenomenal states altogether (and likewise for the intentional capacities Gray et al surveyed).

⁴ These dimensions, according to Gray and Wegner (2009), connect with the distinction between moral *patience* (high on experience, low on agency) and moral *agency* (high on agency, low on experience). For a critique of this work, see Arico (forthcoming).

A clearer statement of the discontinuity thesis we're interested in, and some of the most interesting evidence in favor of it, comes from Knobe and Prinz (2008). Knobe and Prinz acknowledge what we call the continuity view, writing that, "One might have thought that...whatever people turned out to be doing with information about physical constitution, they would at least do the very same thing for all kinds of mental state ascriptions" (70). Ultimately, however, they think that this view is mistaken, concluding that "the process underlying attributions of states that require phenomenal consciousness makes use of information about physical constitution in a way that other mental ascriptions do not" (71). Specifically, Knobe and Prinz contend that people think that group agents have intentional states but not phenomenal states. And they support this contention with evidence on whether people find sentences that *ascribe* various mental states to groups as "sounding natural."⁵

In one of their studies, Knobe and Prinz gave participants sentences that ascribed either a phenomenal state (e.g., "Acme Corp. is experiencing a sudden urge to pursue internet advertising") or an intentional state (e.g., "Acme Corp. believes that its profit margin will soon increase") to a group. Knobe and Prinz then asked participants to rate the sentences on a scale from 'sounds weird' to 'sounds natural.' Participants in this study rated group phenomenal state ascriptions as sounding significantly weirder than group intentional state ascriptions. Knobe and Prinz surmise that the results of this study "seem to indicate that people are unwilling to ascribe to group agents states that require phenomenal consciousness" (75). And if people think groups

⁵ We use the terms "ascription" and "attribution" in importantly different ways. For our purposes, an ascription is a sentence that assigns a state to an entity. An attribution is a mental act of assigning a state to an entity. This highlights the distinction between positing states to an entity at the linguistic level versus the psychological level.

have intentional states but not phenomenal states, that would seem to constitute very good evidence in favor of the discontinuity thesis.⁶

Notice that in this study (as well as other of the studies they discuss) Knobe and Prinz are analyzing uses of and naturalness ratings for various sentences.⁷ Yet their conclusion is not restricted to norms of sentence use; rather, it is about how humans cognize with respect to group mentality. Knobe and Prinz conclude that people are “realists” about group intentionality—that people think that groups actually have intentional states (though not phenomenal states) above and beyond the states of their members. But is this the best interpretation of the data? How, we wonder, can knowing how “natural sounding” people find sentences attributing beliefs and desires to groups shed light on whether people really *think* groups have beliefs and desires?

By analogy, suppose that (as seems likely) people find it natural to say, “Manning was a bundle of nerves.” Clearly, we should not thereby conclude the folk are intuitive proponents of a neurological identity theory. The example illustrates that an assessment of what sentences people find appropriate supports a theory of what people really think only on the assumption that the relevant sentences are being interpreted in a way consonant with the theory. To draw conclusions about how people think about the mental states of groups from naturalness ratings for sentences about group mental states, Knobe and Prinz would have to assume that people are assigning a realist interpretation to the ascriptions—that people interpret the group mental state ascriptions

⁶ Both Arico (2010) and Sytsma and Machery (2009) challenge Knobe and Prinz's conclusions based on the failure to utilize minimal pairs in their stimuli. Arico and Sytsma and Machery also present data suggesting that the original difference between ratings for intentional and phenomenal attributions observed by Knobe and Prinz essentially vanishes once the stimuli are balanced to include matching amounts of information.

⁷ In one notable exception, Knobe and Prinz support their claim using results from the Google search engine. Knobe and Prinz entered intentional attributions to Microsoft (e.g., “Microsoft intends” and “Microsoft believes”) and phenomenal attributions to Microsoft (e.g. “Microsoft feels depressed” and “Microsoft experiences joy”) into Google. Google returned significantly more results for intentional attributions than for phenomenal attributions, which Knobe and Prinz cite as evidence of a cognitive resistance to attributing phenomenal states to groups. However, in an unpublished manuscript, Arico argues that these results are evidence, not of a cognitive resistance to group phenomenal states, but of a general imbalance of online language use. He reports an identical asymmetry in Google results for intentional and phenomenal attributions to individuals (e.g., President Bush and Bill Gates), as well as an asymmetry in results for intentional and phenomenal language in general (i.e., without any subject).

they find ‘natural sounding’ or ‘weird sounding’ as *actually* attributing intentional and phenomenal states to groups over and above their members. But is this assumption warranted?

Knobe and Prinz (2008) defend realism about group mental state ascriptions, contending that, “most researchers who have thought seriously about these questions...think that the expressions under discussion here are not just shorthand and that people really are ascribing mental states to groups” (72). Building on the work of these researchers, Knobe and Prinz seem to offer two main arguments in favor of realism.

One of Knobe and Prinz’s arguments (which they attribute to Huebner, 2008) exploits certain similarities between ordinary ascriptions of mental states to groups, on the one hand, and to persons, on the other. Specifically, they point out that many people think patterns of electro-chemical activity across complex systems of individual neurons ultimately determine a person’s behavior. Nonetheless, people who think this commonly explain human behavior not in terms of neural states, but in terms of abstract psychological generalizations. What’s more, such generalizations are ‘robust’; “they would continue to hold even if the properties of the individual neurons had been somewhat different” (72). In these respects our ascriptions of mental states to groups are incredibly similar to our ascriptions of mental states to persons. People think patterns of human activity, often across complex systems of institutional hierarchy, ultimately determine a group’s behavior. Nonetheless, people also use abstract psychological generalizations to explain the behavior of groups. And these generalizations are robust in just the way generalizations about the behavior of persons are; they would continue to hold even had there been different individual members of the group. Knobe and Prinz argue that given the fact that people are committed to realism about human mental states, and given that these so closely

resemble our ascriptions of mental states to groups, we should also be committed to realism about group mental state attributions.

In light of these similarities between the ways in which people ascribe mental states to groups and persons we may lose sight of numerous dissimilarities. People speak of groups but not persons merging and disbanding, being formed or created whole-cloth, and being coincident with other groups (i.e., composed of the same, numerically identical parts). But a deeper problem with this argument is that it proves too much. After all, people think patterns of electro-chemical activity across complex systems of individual components ultimately determine a car's behavior. Nonetheless, they commonly explain car behavior in terms of abstract psychological generalizations. And these generalizations are robust; they would hold even if the properties of individual components had been somewhat different. Are we then committed to psycho-automotive realism? Of course not. Thus the argument overgeneralizes.

Knobe and Prinz also seem to suggest a second argument in favor of realism about group mental state ascriptions. They explicitly recount two arguments in favor of group minds—the metaphysical view that groups really have certain mental states over and above the individual mental states of their members. These arguments are drawn from Velleman (1997) and Pettit (2003). According to Velleman's argument, you can only intend to do something if you believe you can do it, but (for example) no member of a philosophy department believes she can hire a job candidate, therefore, the intention to hire must be a full-fledged group intention. According to Pettit's argument, a "majority rules" account of group belief would lead to instances of incoherent group belief sets. So, ascribing beliefs to groups requires something more than the majority rules method, something like robust group beliefs. Importantly, these are both arguments in favor of a metaphysical claim; they are *not* arguments explicitly about folk

psychology. Velleman and Pettit are arguing that there are full-fledged group mental states, but their arguments lend no direct support to the descriptive claim of realism about group mental state ascriptions.

Perhaps Knobe and Prinz are tacitly assuming that the most straightforward way to interpret group mind talk is as actually attributing mental states to groups. So, if there is no special burden in accepting realism about group mind talk, we should accept it. If this is the argument Knobe and Prinz mean to be making, it rests on a bad inference pattern. The inference pattern of this argument is as follows:

1. S says, “p”.
2. Strictly, *p* is true.
3. The most straightforward way to interpret “p” is *p*.
4. Therefore, S means *p* when S says “p.”

But consider this argument, which conforms to the relevant inference pattern:

1. You say, “Manning was a bundle of nerves.”
2. Strictly, *Manning is a bundle of nerves* is true (let us suppose).
3. The most straightforward way to interpret, “Manning was a bundle of nerves,” is *Manning was a bundle of nerves*.
4. Therefore, you mean *Manning was a bundle of nerves* when you say, “Manning was a bundle of nerves.”

Clearly, this argument fails to establish that you mean *Manning was a bundle of nerves* when you say, “Manning was a bundle of nerves.” Nor can this form of argument establish realism about group mental state ascriptions.

It is a platitude of many theories of communication that people generally understand the thought expressed by a sentence as corresponding to what the sentence says,⁸ but this is not invariably the case. After all, the aforementioned sentence, “Manning was a bundle of nerves,” presumably says *that Manning was a collection of nerve cells*, but in most contexts it would not be understood to express that thought. We easily recognize that, in most contexts in which this sentence would be used, it would be used to express a different thought—one perhaps not a candidate for accurate and succinct verbal expression. We maintain that similar considerations could lead people to understand group mental state sentences as being used to express something that may not correspond to what those sentences actually say. Even ordinary speakers who reject group intentionality may decide to be brief, saying simply that the corporation *believes* that its profit margin will soon increase as a result of new policies, rather than discussing complex internal debates that result in official press releases. In that case, one would be opting for economy over accuracy, for speaking simply over saying what one *really* thinks is happening. These considerations suggest an alternative interpretation of the aforementioned evidence in favor of the discontinuity thesis.

The fact that there is a plausible alternative to the discontinuity account doesn't, of course, show that this alternative is correct. Given the uncertainty about whether people are

⁸ Here we adopt (what Bach 2001 calls) Grice's Syntactic Correlation Constraint, and equate what is said to the lexical-compositional meaning of the pronounced words. Perhaps what the sentence, “Acme Corp. believes that its profit margin will soon increase,” says is not the realist proposition *that Acme Corp. over and above its members believes that its profit margin will soon increase*. Ultimately, the semantic content of this sentence (and related sentences) is to be determined by the semanticist. Nonetheless, we grant this assumption to the discontinuity theorist, with some support from the semanticist. Link (1984) and Landman (1989 a & b) discuss in detail why we should introduce groups, understood as “plural individuals that are distinct from sums of singular individuals” (Landman 572) into our semantics. For one simple argument, suppose that Acme Corporation consists entirely of Biff, Max, and Sal, all of whom are masters of finance. Then we could say, “Biff, Max, and Sal are masters of finance.” And if group names referred only to sums of singular individuals, we should also be able to say, “Acme Corporation are masters of finance.” But (as Landman notes) constructions such as this seem awkward and perhaps ungrammatical. We want, instead, to say, “Acme Corporation consists of masters of finance.” But we can do that only if we assume that “Acme Corp” refers to, “an individual that does not stand in the part of relation to the sum of the...[masters of finance]...but in a different, consists of, relation” (Landman, 572).

interpreting what is attributed by the sentences in a realist fashion, it would be nice if we could devise a method to get at how people really understand group mental state assertions. Do they really mean to be attributing mental states to groups when they use and endorse these sentences, or are they instead subtly referring to the distributed states of group members and the complex internal policies of groups with these verbal ascriptions?⁹ We think the latter most accurately captures what people have in mind, and in the next section, we present some evidence in favor of that hypothesis.

2. Who Do You Think Thinks When Boeing Thinks?

To get clearer insight into how people are interpreting intentional and phenomenal group state ascriptions, we need an experimental paradigm that neither places participants in a position where a response would force a clash between economy and accuracy, nor puts them in a situation where the apprehension of a potential clash might color their responses. With this goal in mind, we designed a pair of studies based on a paraphrase method previously developed by Phelan (2010).

We first conducted a very informal pilot study in which people were asked to paraphrase group mental state sentences, such as, “Microsoft wants to increase profits”. We thought that people might give a distinctive rendering of “wants”, one that would blunt a realist interpretation. This was not the case at all – every single paraphrase used “want” or a closely related term, like

⁹ These are the two possibilities endorsed in the literature. For instance, some philosophers offer ordinary language arguments that take everyday sentences as literally attributing mental states to groups. Margaret Gilbert (1996), for instance, has argued that our everyday concept of social groups takes them to be “a special kind of thing, a ‘synthesis *sui generis*’”(268), capable of their own subjective states, of holding beliefs that none of the members hold themselves. Deborah Tollefsen (2002) defends a similar view. Others, however, argue that attributions of collective intentionality ought to be analyzed in terms of the mental states of the groups’ individual members. John Searle (1990, 1995), for example, denies that our metaphysics allows for any consciousness beyond individual consciousness, and so all talk about intentionality, including collective intentionality, must be based in an analysis of individuals as the bearers of intentionality. Michael Bratman (1993, 1999) argues that shared intentions are complexes of individual intentions and plans, all interrelated to each other.

“desire”. Something interesting did turn up, however. When people paraphrased these sentences, they tended to use plural pronouns to refer to the group. Obviously, this is quite different from how anyone would paraphrase a sentence attributing an intentional state to an individual. Suppose, for instance, that you were asked to paraphrase the sentence, “Mary believes p .” Though you probably think that Mary’s believing p depends upon the state of her neurons, you would never paraphrase this sentence by saying, “they believe p ,” meaning to refer to Mary’s neurons. This is presumably because you think the belief is fully *Mary’s belief*, regardless of whatever supervenience relation it may bear to her neuronal states.

The difference between how our pilot subjects paraphrased sentences ascribing mental states to *groups* and how we would intuitively paraphrase sentences ascribing mental states to persons raises the possibility that people aren’t really attributing intentional states to corporations when they judge intentional state ascriptions to sound natural. However, the difference could also simply be due to some systematic distinction in the way people paraphrase utterances about groups and persons. Perhaps people quite generally use plural pronouns for groups. Sentences about other states, such as Boeing’s loss of market share, provide a nice contrast case. Such states belong to Boeing rather than to the individuals on whose actions such non-mental states nonetheless supervene.

Thus, in our initial study, we decided to investigate whether there are systematic differences in the pronouns people choose to replace group nominalizations for various kinds of sentences. The possibility that our pilot results were merely the consequence of a general tendency to use plural pronouns for groups predicted that people would utilize plural pronouns for all kinds of group attributions. If, on the other hand, the pronoun usage in our pilot was reflecting something about how people interpret group mental state sentences, then the continuity

view that we are endorsing predicted that people would prefer plural pronouns for both kinds of mental state attribution, but not the non-mental attributions. Knobe and Prinz (and the discontinuity thesis) seemed to predict that people would utilize a plural pronoun only for phenomenal state attributions.

Study 1: Pronoun Replacement Task

Sixty people spending time on the Yale University campus were recruited to participate in our first study and compensated with candy or the sincere gratitude of the experimenter. Participants averaged 22 years of age and 35 were female (one participant did not supply gender information). 81% of participants were between the ages of 18 and 22 (no one under the age of 18 was allowed to participate). Participants came from a variety of majors and averaged one philosophy course.

Each participant was randomly assigned to one of three conditions and given a brief questionnaire involving several sentences. One questionnaire involved non-mental group ascriptions (e.g. “After Boeing lost the Army contract, Boeing needed to lay off workers”); another involved intentional mental state ascriptions (e.g., “After Boeing lost the Army contract, Boeing expected to have to lay off workers”); and a third involved phenomenal mental state ascriptions (e.g., “After Boeing lost the Army contract, Boeing felt anxious about having to lay off workers”). Each condition included four test sentences (i.e. group ascriptions) and four filler sentences. For each sentence, participants were asked to replace the second, underlined instance of the group name (or noun phrase) with either “it” or “they”.¹⁰ For each participant we calculated a distributivist score by awarding 1 point each time the participant selected “they” to

¹⁰ Versions of all test materials are available online at: <https://viewer.zoho.com/docs/gxm5j>

replace the underlined name in one of our test sentences, and 0 points each time the participant selected “it”. Given that we had four test sentences, a distributivist score of 4 was indicative of extreme distributivist sentiments, while a score of 0 indicates extreme realism about the particular kind of group state ascriptions.

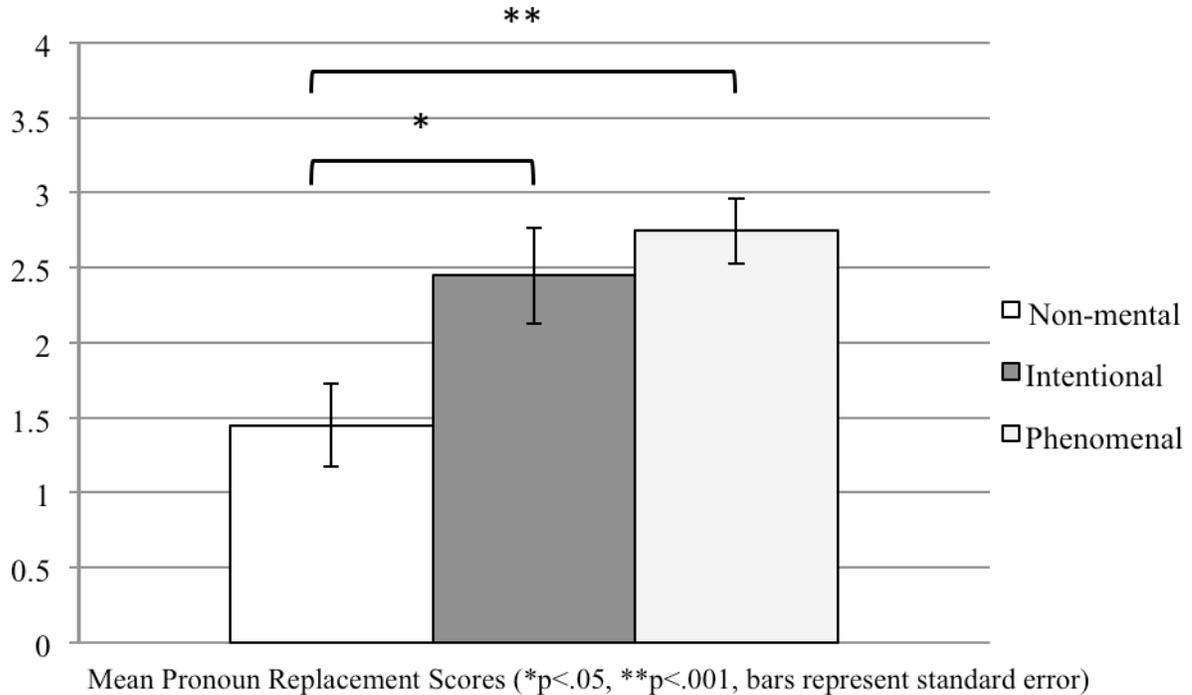
Our prediction was that participants would be more inclined to use the “it” pronoun in the condition involving non-mental state ascriptions than in either mental state condition. That is, we suspected that people were distributivists about both kinds of group mental state sentences (i.e., understanding the attribution as shorthand for distributing the attribution to the members—or some relevant subset of members—composing the group), but realists about ascriptions of other, non-mental states to groups. This hypothesis was borne out. Assuming an alpha level of 0.05 (as we do for this and all other statistical tests reported in this paper), participants were significantly more likely to select the plural pronoun paraphrase for intentional mental state ascriptions (mean distributivist score: 2.45) than they were for non-mental state ascriptions ($m = 1.45$).¹¹ They were also significantly more likely to select the plural paraphrase for phenomenal mental state ascriptions ($m = 2.75$) than for non-mental state ascriptions.¹² However, no significant difference emerged between scores for intentional and phenomenal ascriptions.¹³ Sentences were counter-balanced, and no order effects emerged.¹⁴

¹¹ Results for pair-wise comparisons were subjected to T-Tests. For intentional v non-mental: $t(40) = 2.36, p = 0.023$ (two-tailed), SD (intentional) 1.43, SD (non-mental) 1.23, *Cohen’s d* = 0.75

¹² $t(40) = 3.7, p < 0.001$ (two-tailed), SD (phenomenal) 0.967, *Cohen’s d* = 1.17

¹³ $t(40) = 0.777, p = 0.422$ (two-tailed), *Cohen’s d* = 0.24

¹⁴ An anonymous reviewer raises an important worry for our stimuli: “the distributive question is... potentially affected by whether the group is simply a collective or a more complexly organized unit.” For example, Boeing is widely known to be a multinational corporation, with complex operations, and a variety of employees with vastly different positions ranging from CEO to salesperson to engineer to janitor. The secret task force, on the other hand, is presumably a simple group consisting of only a few members relatively equal to one another in power and influence. It is unreasonable to simply assume that these vastly different kinds of groups will all be thought of in the same way. Perhaps the less complexly organized groups are driving the distributivism effect. To respond to this worry, we ran a series of paired sample T-tests comparing responses for each group to every other group across each of our three conditions. Three significant effects emerged. Distributivism ratings for Phi Lamb ($M = 0.75$) were significantly higher than for Boeing ($M = 0.5, p = 0.021$) in the intentional state condition. Distributivism ratings



It is also useful to consider median results, since these are less susceptible to the effect of outliers than are mean scores. The median distributivist scores for our initial study were:

Non-mental: 1

Intentional: 3

Phenomenal: 2.5¹⁵

Again, these scores suggest that the way people understand both varieties of group mental state ascriptions is quite different from the way they understand ascriptions of mental states to persons.

for Phi Lamb (M = 0.85) were also significantly higher than for Boeing (M = 0.55, p = 0.03) in the phenomenal state condition. However, distributivism ratings for Phi Lamb were significantly higher than for the Secret Task Force (M = 0.55, p = 0.03) in the phenomenal state condition. The difference between MADD (M = 0.8) and Boeing was also approaching significance (p = 0.56) in the phenomenal state condition. These results do not clearly support the contention that mental states are more realistically ascribed to larger, more complex groups than to smaller ones. After all, Phi Lamb is a national fraternity, and even a local chapter is presumably much more complex and hierarchical than a task force. MADD is a national organization, if not as complex as Boeing. Nonetheless, the relatively low distributivism ratings for Boeing do point to the need for careful future work on this question.

¹⁵ The median score here is 2.5 because there were an even number of subjects, and the two scores in the middle of the range were 2 and 3.

Importantly, our method does not rest on assumptions about whether people are interpreting the sentences in a realist or distributivist manner. Instead we infer that people are interpreting group mental state ascriptions distributively, as this affords the best explanation of the significant differences that exist between non-mental state ascriptions, on the one hand, and both varieties of mental state ascriptions, on the other. Furthermore, participants' responses in this study are not subject to concerns of effort. Participants are choosing between equally effortful acts of ticking the it-box or the they-box. These responses therefore provide some of the first evidence about how people are interpreting group mental state ascriptions and, thereby, offer a new perspective on ordinary beliefs about group minds and, by extension, the discontinuity view. They suggest that the impression of discontinuity given by previous studies is misleading.

Study 2: Sentence Replacement Task

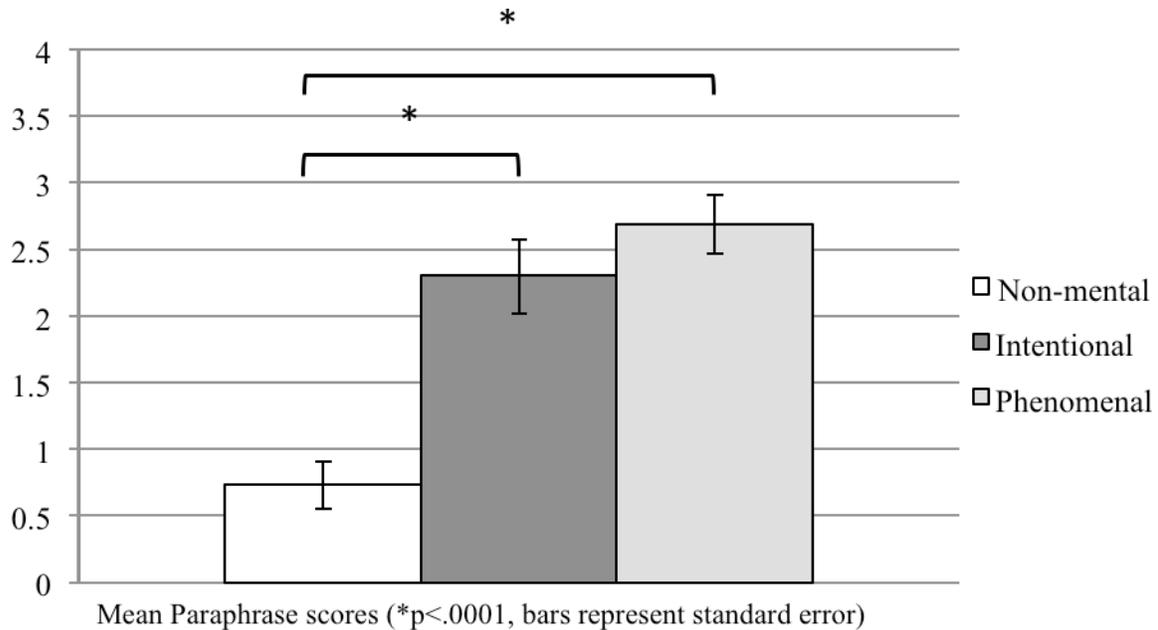
Although study 1 clearly does not rely on the assumption that people are assigning realist interpretations to the test sentences, one may nonetheless worry that it is not altogether free of interpretive assumptions. After all, plural pronouns can also be used non-distributively. 'They' can sometimes be used in American English as a substitute for a singular pronoun, when, for example, the referent's sex is unknown or indefinite. Thus, one often hears things like "A person should always trust their first instinct." 'They' can also be used to refer to a collective (rather than to the members that constitute it), as when one says of a truck-load of shampoo bottles, "They weigh a ton." Perhaps 'they' is being used in the context of our study in some such singular-substitutive way. Of course, such a challenge to our findings would not by itself explain the asymmetry between mental and non-mental ascriptions that we found in our first study. Nonetheless, to support our interpretation of how people are using the plural pronoun 'they' in

the context of our first study, we conducted a second study that attempted a more direct examination of the ordinary understanding of group mental state ascriptions.

Eighty-three students at the University of Arizona participated in this study. Students were recruited from introductory philosophy courses. Participants were once again assigned to one of the three conditions described in Study 1 and given a brief questionnaire involving four sentences. Participants received the same non-mental, intentional, and phenomenal sentences described above. However, rather than being asked to replace noun phrases with pronouns, participants were given two sentences and asked to select which one best described what was meant by the original target sentence. One of these paraphrases was *distributivist*, in that it described the ascription as being directed at some relevant subset of members in the group; and one of the paraphrases was *realist*, in that it described the ascription as being meant for the group itself, not just its members. For instance, participants in the phenomenal condition were asked to select which of the following best described what was meant by “the Secret Task Force felt threatened by the public reaction”: (i) *The majority of the relevant task force participants felt threatened by public reaction*; or (ii) *The task force itself, not just the participants, felt threatened by public reaction*. As before, 1 point was assigned for each distributivist answer, 0 for each realist answer.

Again, our prediction was that participants in the mental (intentional and phenomenal) conditions would offer significantly higher distributivist scores than those in the non-mental condition. And again, the results bore out this prediction. Participants in both the phenomenal and intentional conditions scored significantly higher ($m = 2.69$ and 2.30 , respectively) than participants in the non-mental condition ($m = .73$). *T*-tests revealed significant differences both

between phenomenal and non-mental conditions¹⁶ and between intentional and non-mental conditions.¹⁷ However, responses to phenomenal and intentional conditions were not significantly different.¹⁸



The same trend is evident when we consider the median scores:

Non-mental: 0

Intentional: 2.5

Phenomenal: 3

The data suggest that, contrary to the deflationary explanation sketched above, subjects in our first study likely were *not* using ‘they’ as a singular-substitutive pronoun. Rather, people seem to be thinking about group mental state ascriptions in a largely distributive way. Together, the two studies constitute evidence in favor of the view that people understand these verbal ascriptions of mental states in terms of the mental states of group members.

¹⁶ $t(53) = 6.743, p < .0001, SD(\text{phenomenal}) 1.198, SD(\text{non-mental}) .92, \text{Cohen's } d = 1.85$

¹⁷ $t(51) = 4.706, p < .0001, SD(\text{intentional}) 1.44, \text{Cohen's } d = 1.32$

¹⁸ $t(54) = 1.116, p = .269, \text{Cohen's } d = 0.30$

Of course, one may be hesitant to adopt our distributivist account of the relevant group state ascriptions. For one may wonder how, if distributivism about group mental state ascriptions were correct, we could explain the apparent asymmetry Knobe and Prinz observed in naturalness judgments for intentional vs. phenomenal states. There is, after all, something intuitive and compelling in the report that people judge ‘Microsoft is feeling excruciating pain’ to sound less natural than ‘Microsoft intends to release a new product.’ A failure to explain such asymmetries may outweigh evidence in favor of distributivism. How, then, can we accommodate the intuitiveness of Knobe and Prinz’s results while maintaining the distributivist reading of folk interpretations? In the next section we will offer an explanation that appeals to asymmetries in the roles that group members are saliently thought to occupy.

3. The Qua Members Principle

It is hard to resist the intuition that, ‘Microsoft believes certain things,’ sounds more natural than, ‘Microsoft feels sad about certain things.’ How can we explain this, assuming that both utterances are understood to encode attributions of states to human members, equally able both to believe and feel sad? We maintain that ascriptions of mental states to groups are understood distributively, in terms of attributions of mental states to the individuals that make up groups. However, as we will argue, the relevant individuals over which such attributions range are not always understood as people *simpliciter*. Instead, group mental state ascriptions are understood in terms of attributions of mental states to group members *qua members* of the relevant groups. If we are correct, asymmetries in appropriateness ratings for group state ascriptions are akin to asymmetries for ascriptions of a single state to a particular individual understood as the occupant of different roles.

Consider Bob, employee of Microsoft, inveterate sports fanatic, and father to Jimmy, the high school baseball team's shortstop. It seems more appropriate to say that, "As a father, Bob wants Jimmy to play harder than the other kids on the team," than to say, "As an employee of Microsoft, Bob wants Jimmy to play harder than the other kids on the team." What can explain this asymmetry between mental state ascriptions to a single individual? Presumably Bob wants Jimmy to play harder than the other kids whether we think of him as occupying the role of father or Microsoft employee. But the asymmetry seems mysterious only because we are now thinking of Bob as a person simpliciter, to whom this state is always appropriately attributed, not as the occupant of a role, as the previous ascriptions invited us to think of him. It is qua occupant of the role of Jimmy's father that Bob wants Jimmy to play harder than the other kids. To think of Bob simply as a person obscures our conception of Bob as an occupant of the roles for which our previous attributions were appropriate or inappropriate.

But what is it to think of a person as an occupant of a role? We propose that a particular state is appropriately attributed to a particular person qua occupant of a particular role, insofar as the person actually occupies the state, and the state is saliently associated with the relevant role.¹⁹ Verbal ascriptions of a state to a person qua occupant of a role will be deemed appropriate insofar as the states ascribed are appropriately attributed to the person qua occupant of the role according to our proposal. Our proposal, in turn, is supported in so far as it conforms with patterns of appropriateness judgments relative to qua-occupant ascriptions.

It is often natural to understand and interpret attributions to individuals through the lens of their membership in groups. For instance, intuitively, it would be more appropriate to say that,

¹⁹ There are a variety of ways a state may come to be saliently associated with a particular role. For example, we may associate a given state with a given role because occupying the state is morally normative, statistically normative, or stereotypically associated with the role. Furthermore, salient associations are relative to individual beliefs (though we expect high levels of commonality in such beliefs, as with prototypicality judgments) and context sensitive.

“As a member of Al Qaeda, Bob hates America,” than to say that, “As a father, Bob hates America.” Hating America is saliently associated with being a member of Al Qaeda, thus one supposes that (in the absence of other defeating facts about Bob’s biography) hatred for America is a critical factor in his membership in Al Qaeda. It is less obvious, however, how Bob’s being a father would relate to his hating America. Thus, *qua Al Qaeda member*, Bob’s hating America is intuitively plausible in a way that, *qua father*, it is not.

We are now in a position to explain asymmetries between group ascriptions of intentional and phenomenal states. We contend that asymmetric attributions of mental states to groups are explained in virtue of a general Qua Members Principle:

QM Principle: When a state is not appropriately attributed to a group over and above its members, verbal ascriptions of the state to a group are understood distributively in terms of attributions of states to (a set of) group members *qua members* of the relevant group, where states are appropriately attributed to the group members *qua* occupants of the role of group member insofar as the people who make up the group are thought to occupy the relevant states and the states are saliently associated with the role of being a member of the group.

If the QM Principle is correct, the asymmetries for group mental state *ascriptions* are due to asymmetries in state *attribution* appropriateness to group members. For example, it strikes us as more appropriate to say that, “Al Qaeda hates America,” than to say that, “Microsoft hates America,” because the state of hating America is saliently associated with the members of Al Qaeda’s role *as Al Qaeda members*. Analogously, the asymmetries between intentional and phenomenal state ascriptions that Knobe and Prinz observed—along with the intuitively asymmetric ascriptions mentioned at the beginning of this section—can be explained as the result of asymmetries between how appropriate it is to attribute various intentional and phenomenal mental states to corporate employees. Experiencing great joy, feeling excruciating pain, and experiencing a sudden urge to pursue internet advertising, are less saliently associated with the job description of a corporate employee than are believing the profit margin will

increase, intending to release a new product, or deciding to adopt a new marketing plan. Thus, we propose that Knobe and Prinz observed their asymmetry between intentional and phenomenal state ascriptions in virtue of the fact that they asked how natural ascriptions of these states were for a generic corporate entity. Put the right group into the ascriptions, one for whom certain phenomenal states are saliently associated with the role of being a member, and we should find that phenomenal states are just as likely to be attributed to groups as are intentional states.

Study 3: Testing the QM Principle

Can the QM Principle predict ordinary appropriateness judgments for group state ascriptions? Notice that according to the QM Principle, the states a group member will be thought to occupy are the intersection of states the person simpliciter occupies and the states that are saliently associated with the role of being a member of the particular group.²⁰ Thus, we can test the QM Principle by holding one aspect of this equation fixed while manipulating the other. This is precisely what we did in our study. We devised a test in which participants were constrained to attribute the same state to a large proportion of the people who make up two different groups, but in which the states so attributed are more saliently associated with being a member of one group or the other. Thus, for each state, though the *people* who make up each group occupy the relevant state, it is more appropriate to attribute the state to *members* of one group or the other; and in virtue of this asymmetry for group member attributions, we hypothesized a similar asymmetry in how appropriate people would find ascriptions of the relevant states to the relevant groups.

Obviously, we are interested in demonstrating the role of the QM Principle for phenomenal and intentional group mental state ascriptions. However, there is nothing special about mental state ascriptions so far as the QM Principle is concerned. The principle kicks in

²⁰ In point of fact, this is an over-simplification, since salient association clearly comes in degrees.

whenever a state is not appropriately attributed to a group over and above its members. In our previous studies on the distributivism of group mental state ascriptions, the non-mental states we investigated—such as losing donations or meeting with public condemnation—were states that could be appropriately attributed to groups over and above their members. But this is not true of all non-mental states. A group cannot dress well or get drunk over and above its members. Thus, non-mental state ascriptions involving these and other non-mental states should be understood distributively and thought appropriate according to the precepts of the QM Principle. In other words, the difference between non-mental and mental group state ascriptions is that some non-mental states are appropriately attributed to groups over and above their members, whereas no mental states are. Thus, we devised a 2 (association: associated, not associated) x 3 (state: non-mental, intentional, phenomenal) study to investigate the QM Principle.

Two hundred and eighteen participants were recruited from Amazon's Mechanical Turk and compensated with \$0.30 or \$0.35. Participants averaged 38 years of age and 140 (64.2%) were female. Each participant was assigned to one of six conditions. In each condition, participants made an appropriateness assessment for each of two state ascriptions. For each state, participants read a vignette describing a target group and stipulating that a significant proportion of the people who make up the group came to occupy a particular state. They were then asked to consider the appropriateness of a target group state ascription and to make their assessment on a five-point Likert scale (anchored at, "inappropriate," "in between," and "appropriate"). So, for example, participants in the associated intentional state condition read the following story:

The American Red Cross is a volunteer-led, humanitarian organization that provides emergency assistance. On August 28th of this year, members of the Red Cross learned that a powerful hurricane had devastated portions of the east coast of the United States.

And assessed the following ascription in terms of appropriateness:

The American Red Cross came to believe that a hurricane had hit the east coast.

Participants in the unassociated intentional state condition read the following story:

Microsoft is a public multinational corporation that develops and manufactures a wide range of computer products and services. On August 28th of this year, employees of Microsoft learned that a powerful hurricane had devastated portions of the east coast of the United States.

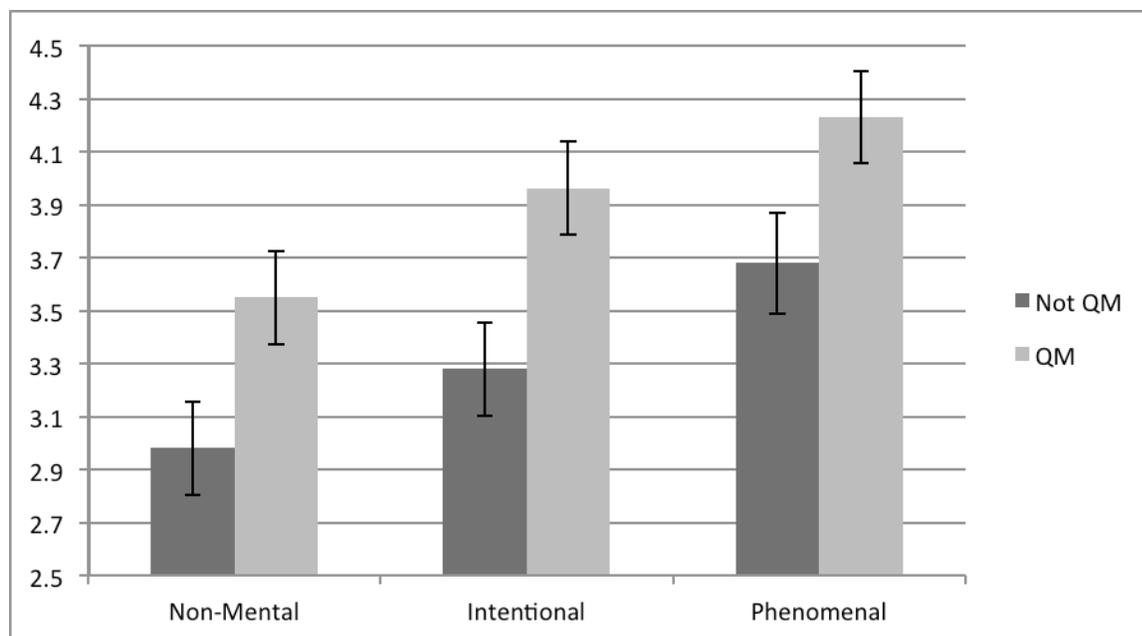
And assessed the following ascription in terms of appropriateness:

Microsoft came to believe that a hurricane had hit the east coast.

Each participant's ratings for the two ascriptions within each condition were then averaged, to generate an overall appropriateness rating.

The QM Principle predicts that participants will find group state ascriptions appropriate insofar as the relevant state is typically associated with being a member of the relevant group. The results of our third study bore out this prediction. Across each level of the state factor (non-mental, intentional, and phenomenal), participants deemed group state ascriptions more appropriate when the specific state was associated with being a member of the target group.²¹

²¹ Means and standard deviations for non-mental: Not saliently associated ($M = 2.94$, $SD = 1.13$), saliently associated ($M = 3.55$, $SD = 1.13$). Means and standard deviations for intentional: Not saliently associated ($M = 3.26$, $SD = 1.24$), saliently associated ($M = 3.96$, $SD = 0.8$). Means and standard deviations for phenomenal: Not saliently associated ($M = 3.63$, $SD = 1.18$), saliently associated ($M = 4.24$, $SD = 0.87$).



A two-way between-subjects analysis of variance was conducted to evaluate the effect of salience and (kind of) state on participants' overall appropriateness ratings. We found a main effect for salience,²² and a main effect for kind of state.²³ However, there was no interaction effect.²⁴ Furthermore, post hoc comparisons using the Fisher LSD test revealed that participants were more happy with group ascriptions of the target phenomenal states (being happy, being sad) than the target intentional states (believing, wanting) or non-mental states (attending, getting drunk). Participants in the phenomenal conditions gave significantly higher overall appropriateness ratings than those in the intentional conditions ($p < .05$) or the non-mental conditions ($p < .001$). And participants in the intentional conditions gave significantly higher overall appropriateness ratings than those in the non-mental conditions ($p = .04$).

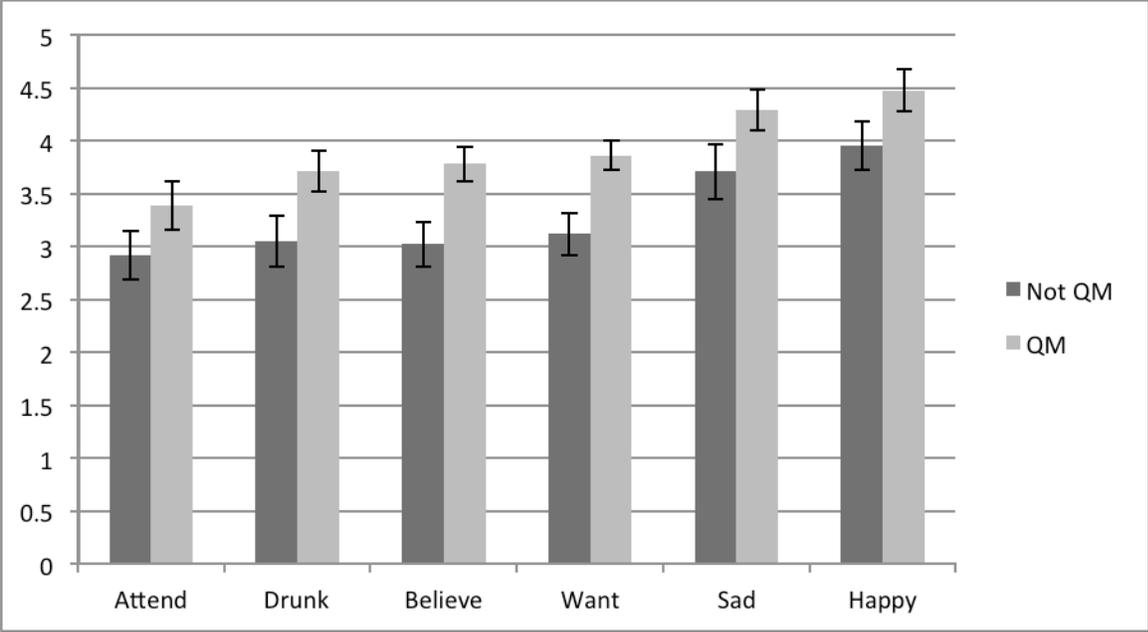
We can gain more insight into the QM Principle by setting aside the overall appropriateness ratings for the three different kinds of states and looking instead at the pattern of results for each specific state. For each state we find a pattern consistent with the QM Principle.

²² $F(1, 212) = 19.535, p < .001$

²³ $F(2, 212) = 7.51, p = .001$

²⁴ $F(2, 212) = .046, p = .955$

Regardless of the state, group ascriptions were always deemed more appropriate when occupying the state is typically associated with being a member of the specific group under consideration.



Results for each state were subjected to a T-test comparing the associated condition to the not associated condition. For five of the six states we found a significant difference in the direction predicted by the QM Principle, with a non-significant trend in the appropriate direction for *attending*.²⁵

One may have found it prudent to reject distributivism even in light of the results of our paraphrase studies if one felt that it would leave us without any plausible explanation of asymmetries in group state ascriptions. But the distributive QM Principle has no trouble explaining the asymmetries. Thus, the QM principle conforms best to the totality of the evidence. Verbal ascriptions of group mental states are understood distributively in terms of attributions of states to group members qua members of the relevant group, and they are deemed appropriate

²⁵ See <https://viewer.zoho.com/docs/gxm5j> for detailed statistics

insofar as the states are saliently associated with the role of being a member of the relevant group.²⁶

4. Impact of the Empirical Results

Our discussion thus far has implications both specific to the topic of folk metaphysics of mind and general to the methodology employed in the study of folk psychology. First, our studies suggest that people's attributions of mental states to groups do not align with a realist interpretation of their verbal ascriptions of mental states to groups. But as we have previously discussed, a realist interpretation of group mental state ascriptions has been essential to one line of argument for the discontinuity thesis. Thus, our studies undercut important evidence in favor of a critical discontinuity. If asymmetries in group ascriptions do not support a difference between how the folk attribute intentional versus phenomenal states to groups, then there is no evidence from group ascriptions that folk psychology is sensitive to divergent sets of properties for intuitive attributions of intentional versus phenomenal states.

But the results fit perfectly with what would be predicted by the continuity view. Notice that, in both paraphrase studies, there was no statistical difference between subjects' treatment of intentional ascriptions and their treatment of phenomenal ascriptions. That is, participants did not judge either intentional states or phenomenal states as actually being attributed to groups over and above their members. Rather, they strongly preferred to replace group names with plural pronouns and to utilize distributivist paraphrases for both kinds of mental attribution. In short, they treated the two kinds of group mental state attributions the same.

We turn now to the broader implications of this discussion for the design of studies that try to answer questions about commonsense psychology by investigating how people judge the

²⁶ Obviously we also need an account of how many and what members of a group must be thought to possess a state for it to be appropriate to ascribe the state to the group. We set such considerations aside for future work.

naturalness of mental state ascriptions. We have just argued that people tend to understand verbal ascriptions of mental states to groups in terms of the mental states of group members. Knobe and Prinz's studies don't support a critical discontinuity in the folk metaphysics of mind, because they rest on an unfounded assumption, that people are interpreting group ascriptions realistically. Other theorists have drawn on this same assumption. For example, Huebner, Bruno, and Sarkissian (2010) assess whether participants in China and the United States find that various group ascriptions sound natural (or weird) and conclude "that commonsense psychology in East Asia does not generate nearly as great a gulf between the acceptability of an ascription of a mental state to an individual and the ascription of a mental state to a group as we find in the commonsense psychology of the West" (241). This experimental paradigm, like that of Knobe and Prinz, takes for granted that participants are interpreting these sentences as realist group attributions.

More generally, it is common for researchers investigating a broad array of topics in folk psychology to assume authority on how participants understand various kinds of sentences. This is true not only of experimental philosophical research, but also of research from other disciplines, as a prominent example demonstrates.

The psychologists Paul Bloom and Csaba Veres (1999) showed participants videos (modeled after those used in the seminal Heider and Simmel (1944) study) in which groups of shapes exhibited structured movement, akin to the movements that you might see actors make in a chase scene. Participants were significantly more likely to offer intentional descriptions of these videos, compared with controls in which the same videos were played in reverse (and therefore did not exhibit structured movement). Bloom and Veres conclude that "these results suggest that one can observe the dissociation between the notions of 'intentional entity' and

‘object’ even in the simple context of the perception of moving geometrical figures” (B7). However, the problems we have identified in Knobe and Prinz’s argument (2008) obviously raise the possibility of similar problems for other studies such as this one.

In general, we should be wary of taking participants’ words and putting our own ideas into their heads, and we should be especially wary when possible alternate understandings are readily available. It is easy to see how a concern for economy could lead an experimental participant to describe one of Bloom and Veres’ structured films using a strictly inaccurate intentional gloss. For example, one says that the red rectangles “did not let the dots touch the green rectangles,” because, in that condition, the scripted actions of the groups of shapes lend themselves to personification, and the personification provides a brief way of encoding a lot of information.

The personification suggests more economical but less accurate attributions of motion patterns to geometric shapes. When the same video is played in reverse, an intentional gloss would no longer be strictly, but rather grossly, inaccurate. It would no longer be even suggestively useful. In that condition, experimental participants must make do with costly descriptions of motion patterns. For example, to describe the same scene described above but played in reverse, one participant writes, “The five green dots did not come in contact with the blue dots at this time. However the red and blue groups made contact.” Bloom and Veres transcribe only one forward video description and one reverse description, so we cannot observe any general trend. However, the forward description, glossed in intentional terms, is only 85 words long, whereas the reverse description, which necessarily relies more on trajectory talk, is, at 147 words, almost twice as long. This suggests that people may be availing themselves of inaccurate ascriptions for the sake of brevity in describing the forward videos.

5. Concluding Remarks

Before concluding our paper, we would like to consider how best to categorize the kind of inaccurate talk that we have been exploring. Is it correct to say that people speak figuratively when they attribute states to group members by ascribing them to groups? Well, yes and no. There are, it turns out, two rather different senses of “figurative” in play. The inaccurate talk we have been discussing conforms with the concept of figurative common to philosophers and linguists, which encompasses all utterances involving loose use of terms. However, there is also an ordinary language notion of the figurative, as some of the authors of the current paper (Arico and Nichols) discovered when they set out to show, in a previous study, that subjects regarded attributions of mental states to groups as figurative, rather than literal, claims.²⁷ Contrary to the authors’ prediction, participants claimed that attributions of intentional mental states to groups were *literally* true. We now believe that this is because the participants in this previous study were operating according to the ordinary notion of figurative language, rather than the philosopher’s concept of loose use.

Whatever the ordinary notion of figurative language involves, the results from Arico and Nichols and colleagues suggest that it is missing from the sentence “Some corporations want tax cuts,” but not from “Einstein was an egghead.” “Some corporations want tax cuts” is literal from the ordinary point of view; but this does not confirm the variety of mental state realism that is important to cognitive scientists investigating what people think about others minds. Proponents of the discontinuity thesis want to show that people *actually think* that (e.g.) group entities can

²⁷ Participants were first screened for their ability to distinguish literal and figurative uses of terms by having them assess obviously figurative and literal sentences such as “Einstein was an egghead,” and, “George W. Bush is President of the United States” (which was true at the time of the study). Participants were then presented with a series of sentences attributing different mental states to groups, such as “Some corporations want tax cuts,” and asked to rate these on a scale of literalness.

have intentional mental states. But if we suspect that people are departing from encoded word meaning when they say, “Acme Corp believes that its profits will soon increase,” then we cannot justifiably draw this conclusion.²⁸ To see this, consider that in the aforementioned study, people generally judged the sentence, “The stock market collapsed in the 1920’s,” to be literally true. Are we to conclude from this that people *really think* the stock market physically fell in on itself in the 1920’s? Of course not. The sentence is not figurative from an ordinary point of view, but it is loose and non-literal from the philosopher’s perspective.²⁹

As the example demonstrates, the philosopher’s concept has a broader extension than the ordinary language notion. The ordinary notion of the figurative clearly involves something more than loose use—in other work, one of the authors of the current paper (Phelan) has characterized this something more as a distinctive kind of affective content. In any case, in so far as we are drawing conclusions about what people *actually* think from what people say, it will be important to guard against figuration in the philosopher’s sense of loose use. Loose use, but not ordinary language figuration, is what we contend is afoot in ordinary group mental state ascriptions.³⁰

In raising this worry about using mental state talk to investigate folk theories of mind, we are offering a similar cautionary lesson as Cullen (2009), who argues that many experimental philosophers, “appear to have proceeded on the assumption that intuitions can be simply read off from survey responses” (275). As we have noted, however, this error is also found in the work of experimental *psychologists*. Furthermore, while it is clearly incumbent upon experimentalists working in all branches of the academy to familiarize themselves with general approaches to

²⁸ As in section 1 above (see fn 8), here we concessively assume that the relevant sentences actually ascribe mental states to groups over and above their members. Again, this is a matter to be settled by the semanticist.

²⁹ Of course, it may be that “collapse” encodes a distinct lexical use that encompasses what happened to the stock market. But in that case, the point still stands. Where there is semantic ambiguity we have to be careful about which lexical use is in play in drawing conclusions about what is thought from what is said.

³⁰ Though we have characterized group mental state ascriptions as instances of loose use in the body of this paper, officially we want to remain non-committal as to whether such talk constitutes loose use or merely contributes to implicature calculation.

survey methodology (e.g. Schwarz 1996), we mean to be offering a specific lesson for those using verbal reports to investigate mental state attribution: Careful consideration of the pragmatic pressures of accuracy and economy must inform such methods. In the case at hand, experimentalists should be sensitive to the fact that there are multifarious influences on people's statements about the mentality of various non-biological entities. These pressures might make it difficult for people to pronounce on the mentality of robots, ghosts, and multi-national corporations without speaking a sentence that is, by their own lights, strictly inaccurate.

Works Cited:

- Arico, A. forthcoming. "Breaking Out of Moral Typecasting." *Review of Philosophy and Psychology*.
- Arico, A. unpublished manuscript. "The Folk v. Acme Corp."
- Arico, A. 2010. "Folk Psychology, Consciousness, and Context Effects." *Review of Philosophy and Psychology*, 1(3): 371-393.
- Arico, A, Fiala, B., Goldberg, R., and Nichols, S. 2011. "Folk Psychology of Consciousness." *Mind and Language* 26 (3): 327-352.
- Bach, K. 2001. "You Don't Say?" *Synthese* 128: 15--44.
- Block, N. 1995. "On a Confusion about a Function of Consciousness." *Behavioral and Brain Sciences*, 18(2): 227-287.
- Bloom, P. & Veres, C. 1999. "The Perceived Intentionality of Groups." *Cognition* 71: B1-B9.
- Bratman, M. 1993. "Shared Intention." *Ethics* 104: 97-113.
- . 1999. *Faces of Intention*. Cambridge, MA: Cambridge University Press.
- Cullen, S. 2009. "Survey-Driven Romanticism." *Review of Philosophy and Psychology* 1: 275-296.
- Fiala, B., Arico, A., and Nichols, S. 2011. "On the Psychological Origins of Dualism: Dual-process Cognition and the Explanatory Gap." In E. Slingerland & M. Collard (eds.) *Creating Consilience: Issues and Case Studies in the Integration of the Sciences and Humanities*. Oxford University Press, 88-109.
- Gilbert, M. 1996. *On Social Facts*. Princeton, NJ: Princeton University Press.
- Gray, H., Gray, K., and Wegner, D. M. 2007. "Dimensions of mind perception." *Science*, 315, 619.
- Gray, K., Gray, H. M. and Wegner, D. M. 2008. "A replication of the mind survey." Unpublished data. Harvard University.
- Gray, K. and Wegner, D. M. 2009. "Moral typecasting: Divergent perceptions of moral agents and moral patients." *Journal of Personality and Social Psychology*, 96, 505-520.
- Horgan, T. and Tienson, J. 2002. "The Intentionality of Phenomenology and the Phenomenology of Intentionality." In D. Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford, UK: Oxford University Press, 520-33.

- Huebner, B., Bruno, M., and Sarkissian, H. 2010. "What does the nation of China think about phenomenal states?" *Review of Philosophy and Psychology*, 225-243
- Knobe, J. and Prinz, J. 2008. "Intuitions about consciousness: Experimental studies." *Phenomenology and Cognitive Science*, 7, 67–83.
- Kriegel, U. 2003. "Is Intentionality Dependent Upon Consciousness?" *Philosophical Studies* 116 (3): 271-307.
- Landman, F. 1989a. "Groups, I" *Linguistics and Philosophy* 12 (5): 559-605.
- Landman, F. 1989b. "Groups, II" *Linguistics and Philosophy* 12 (6): 723-744.
- Link, G. 1984. "Hydras. On the logic of relative clause constructions with multiple heads" in F. Landman and F. Veltmann (eds.). *Varieties of Formal Semantics*. Dordrecht: Foris
- Nagel, T. 1974. "What is it like to be a bat?" *Philosophical Review* 4:435-50.
- Phelan, M. 2010. "The inadequacy of paraphrase is the dogma of metaphor." *Pacific Philosophical Quarterly* 91: 481-506.
- Robbins, P. and Jack, A. 2006. "The phenomenal stance." *Philosophical Studies*, 127, 1, 59-85.
- Rosenthal, D. 1997. "A Theory of Consciousness" in Block, Flanagan, and Güzeldere, eds. *The Nature of Consciousness*. Cambridge, MA: MIT Press, 729-753.
- Schwarz, N. 1996. *Cognition and communication: Judgmental biases, research methods, and the logic of conversation*. Mahwah: Erlbaum.
- Searle, J. 1990. "Collective Intentions and Actions." In *Intentions in Communication*, P. Cohen, J. Morgan, and M.E. Pollack, eds. Cambridge, MA: Bradford Books, MIT press.
- . 1995. *The Construction of Social Reality*. New York, N.Y.: Free Press.
- Strawson, G. 2005. "Real intentionality v.2: why intentionality entails consciousness" *Synthesis Philosophica* 2 (40):279-297.
- Sytsma, J and Machery, E. 2009. "How to study folk intuitions about consciousness." *Philosophical Psychology*, 22(1), 21-35.
- . 2010. "Two conceptions of subjective experience," *Philosophical Studies*, 151(2): 299-327.
- Tollefsen, D. 2002. "Organizations as true believers," *Journal of Social Philosophy* 33: 395-410