

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

Mark Phelan
Yale University

Adam Arico
University of Arizona

Shaun Nichols
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). In other words, the cognitive process that leads us to attribute mental states to various entities

[†] Versions of this work were presented at the London School of Economics, the Metro Experimental Research Group, the Southern Society for Philosophy and Psychology, the University of Arizona Experimental Philosophy Lab, and the Society of Philosophy and Psychology. Audience comments helped improve this paper. The authors are also grateful to Michael Bruno, Georg Kjöll, Joshua Knobe, Eric Mandelbaum, Justin Sytsma, and Jonathan Weinberg, who gave valuable comments on earlier drafts.

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, 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.

Constructing theories of folk psychology based on how people talk about minds and mental states is a common method within philosophy and psychology.¹ 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 even beyond the group mind debate.

¹ Heider and Simmel (1944), for example, use people’s descriptions of how several geometric shapes move around a movie screen to investigate how motives are attributed to other persons’ behaviors. Experimental paradigms used to assess children’s theory of mind, such as Gopnik and Astington’s (1988) appearance-reality task, rely on children’s claims about the mental states of others. And Bloom and Veres (1999) analyze experimental participants’ descriptions of groups of shapes to assess, “when we...think of groups as single intentional entities” (B3).

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 a key distinction between phenomenal states (such as feeling anxious, getting depressed, and experiencing joy) and non-phenomenal intentional states (such as believing, intending, and wanting) (e.g., Knobe and Prinz 2008; Gray, Gray, and Wegner 2007; Robbins and Jack 2006).²

In this paper, we want to examine some fascinating evidence that people genuinely see a discontinuity between what is required for phenomenally conscious states and what is required for intentional states. We'll call this view the *discontinuity* thesis. According to this thesis, there exists a discontinuity between the physical features that typically dispose people to attribute intentional states to an entity and the physical features that typically dispose people to attribute phenomenal states to an entity. 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

Unlike the continuity view, 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

² 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. Therefore, we will continue to draw the distinction in terms of phenomenal and intentional states in what follows.

³ One example of a continuity view can be found in Arico et al. (2011), where they defend 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', which we are disposed to do whenever we represent that thing as displaying certain relatively simple features (i.e., eyes, distinctive motion, and contingent interactive behavior).

ordinary people distinguish (albeit, perhaps, only tacitly) phenomenal from (non-phenomenal) intentional states. The evidence in favor of this second claim has been contested, both empirically and theoretically, elsewhere (Arico 2010, Sytsma and Machery 2009).⁴ Some have even gone so far as to argue that ordinary people do *not* register the distinction familiar to philosophers and cognitive scientists (e.g. Sytsma and Machery 2010). However, in this paper we remain neutral on this issue. We want to challenge only the primary claim, that there is a discontinuity between the sorts of properties that dispose us to attribute phenomenal states and intentional states.

Now, as a matter of philosophical fact, 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 question of whether there is a genuine discontinuity, it may nonetheless illuminate other philosophically relevant questions.

For instance, a number of philosophers have recently suggested that attributions of conscious states and attributions of intentional states play fundamentally different roles in moral cognition. Knobe and Prinz (2008), for example, present data suggesting that people think of ascriptions of intentional states, like beliefs and memories, as being used primarily “to predict, explain or control behavior” (82), whereas they maintain that “ascriptions of phenomenal consciousness are best understood in terms of their role in facilitating moral judgment” (82). Robbins and Jack (2006) likewise argue that there is an essential connection between seeing a

⁴ 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 also extends this criticism to the studies by Huebner, Bruno, and Sarkissian, 2010). 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.

thing as capable of having phenomenal states and seeing it as an object of moral consideration. As they put it, “regarding something as a locus of experience involves regarding it as a potential target of moral concern... as something that one ought to shield from harm” (70). Lastly, Gray and Wegner (2009) associate the psychological distinction between intentional states and phenomenal states with the distinction between the perception of moral agency and the perception of moral patiency. They argue for a moral typecasting thesis, according to which: (a) our perception of entities as moral *agents* is linked to our attributing those things intentional capacities, such as self-control, memory, planning, and thought; and (b) our perception of moral *patiency* is linked to our attributing a thing states like hunger, pain, pleasure, pride, and joy.

While one may take issue with various details in the above proposals, it remains a live possibility that our moral judgments are tied in radically different ways to our attributions of intentional states and phenomenal consciousness. So, at the very least, 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.

In addition to illuminating the psychology of moral judgment, learning about 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. forthcoming).

2. Evidence for the Discontinuity Thesis

What evidence could tell in favor of the discontinuity thesis? 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".⁵ 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 have one kind of state but not the other. For instance, they maintain that, people regard God as having high levels of Agency but low levels of Experience. In unpublished work, Gray et al (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.

⁵ These dimensions, according to Gray and Wegner (2009), connect with the distinction between moral *patience* and moral *agency*. Moral patients, they conclude, rate high on the experience dimension and low on the agency dimension; moral agents, by contrast, rate high on agency and low on experience.

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) which are simply inappropriate for God and Google to experience. As a result, the evidence on their categories of Agency and Experience doesn’t show that people think some entities have intentional states but no 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).⁶

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 there was some general truth about how people ascribe mental states and 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 traditional 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

⁶ Supporting materials for Gray et al. (2007) discuss the surveyed mental states in more detail and are available online at www.sciencemag.org/cgi/content/full/315/5812/619/DC1

but not phenomenal states. And they support this contention with experimental results from studies intended to survey whether people find sentences that *ascribe* various mental states to groups as “sounding natural.”⁷

Consider one of Knobe and Prinz’s studies, in which they gave participants sentences that ascribed either a phenomenal state or an intentional state to a group. Knobe and Prinz then asked participants to rate the sentences on a scale from ‘sounds weird’ to ‘sounds natural.’ For instance, participants assessed both the intentional sentence, “Acme Corp. believes that its profit margin will soon increase,” and the phenomenal sentence, “Acme Corp. is experiencing a sudden urge to pursue internet advertising.” Participants in this study rated group phenomenal state ascriptions as sounding significantly weirder than group intentional state ascriptions. Indeed, subjects rated the least acceptable group intentional state sentence as more “natural sounding” than the most acceptable group phenomenal state sentence. 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 have intentional states but not phenomenal states, that would seem to constitute very good evidence in favor of the discontinuity thesis.

Notice that in this study (as well as other of the studies they discuss) Knobe and Prinz are analyzing uses of, and appropriateness ratings for, various sentences. 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

⁷ Here we will use the terms “ascription” and “attribution” in importantly different ways. For our purposes, an ascription is a sentence that assigns a mental state to an entity. An attribution is a mental act of assigning a mental state to an entity. In other words, we mean to be highlighting the distinction between positing mental states to an entity at the linguistic level versus the psychological level.

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 thoughts 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 seem to be assuming that people are assigning a realist interpretation to the ascriptions—that they interpret the group mental state ascriptions 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?

Could a theorist support this assumption—that ordinary people interpret group mental state sentences as expressing attributions of intentional and phenomenal states to groups *over and above their members*—by claiming that the sentences themselves actually say (i.e., have as their semantic value) that groups over and above their members have the relevant mental states?⁸

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. It is possible, of course, that 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).

It is normal to suppose, after all, that people understand the thought expressed by a sentence to correspond to what the sentence says. This point is captured by Grice’s (1989) maxim of quality (“Do not say what you believe to be false” (27)) and by contemporary theories of communication. Relevance Theory, for instance, construes interpretation as resting on the assumption that utterances achieve optimal cognitive effects at low processing costs. But the more encoded word meanings need to be adjusted, the higher the processing cost will turn out to be.⁹

We do not think such considerations support a realist understanding of the relevant attributions. Though it is true that people generally understand the thought expressed by a sentence to correspond to what the sentence says, 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 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 attribute a thought that may not correspond to what those sentences actually say. Readily apprehensible conversational pressures may lead participants to reject realist attributions for group mental state sentences.

To see the sort of conversational pressures we have in mind, suppose ordinary people don’t actually think that mental states—either intentional or phenomenal—can be attributed to groups over and above the members of those groups. If people do not think that the behavior of a large corporation like Microsoft, for example, is explained by the corporation’s beliefs and desires (over and above the beliefs and desires of the corporation’s employees), then people

⁹ See Sperber and Wilson, 1995.

presumably think Microsoft's behavior is ultimately explained by some of the employees working in concert, guided by their own individual mental states. In that case, how would one explain, for instance, a corporation's reaction to some opaque new government policies? One might speak of complex internal debates in which relevant employees state their opinions, eventually reaching compromise, leading to a press release detailing the corporation's official position that the corporation's profit margin will soon increase as a result of the new government policies, and so on. Such a verbal explanation certainly would very accurately report one's beliefs about the corporation's behavior. But it would do so at the cost of economy.

One might instead be very brief, saying simply that the corporation *believes* that its profit margin will soon increase as a result of the new policies. In that case, one would be opting for economy over accuracy, for speaking simply over saying what one *really* thinks is happening. Whatever option one chooses, one faces here a conflict between two conversational interests. In Grice's terms, one faces a clash between maxims of quality and of manner (*be brief*). Put in Relevance Theoretic terms, the effort required to process the long statements that transparently reflect what one actually believes would rob the utterance of much relevance—even more than is lost when one says something that isn't entirely accurate.¹⁰ Given that conflict, it would be reasonable for one to utter a statement that does not accurately reflect what one thinks, relying on the audience to interpret the statement about a corporation's beliefs as implicating more accurate (but less economically expressible) thoughts about some vaguely-grasped internal process. These considerations suggest an alternative interpretation of the aforementioned evidence in favor of the discontinuity thesis.

¹⁰ Other contextualists and indexicalists would offer distinct explanations for how and why one pronounces a sentence that on its surface departs from what one means to claim (see Bach, 1994, for a non-relevance-theoretic contextualist account, and Stanley and Szabó, 2000, for an indexicalist approach to the distinct phenomena of quantifier domain restriction). Applied to the cases under discussion, such theories would still acknowledge the same basic conflict between economy and unambiguousness.

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 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 (perhaps) the complex internal policies of groups with these verbal ascriptions?¹¹ 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.

3. 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 conversational maxims 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 small pilot study in which people were asked to paraphrase group mental state sentences, such as, “Microsoft wants to increase profits”. We thought that

¹¹ These seem to be 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. Likewise, Deborah Tollefsen (2002) claims that “in everyday discourse... we often attribute beliefs, desires, and other intentional states to groups. Even the most ardent individualists routinely apply the intentional idiom to groups like corporations, nations, and committees” (26). Others, however, argue that attributions of collective intentionality ought to be analyzed in terms of the mental states of the group's 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. Thus, while no one individual may possess the collective intention, the collective intention is nonetheless captured by a complex distribution of mental states amongst the individual members.

people might give a distinctive rendering of “wants”, one that avoided realist implications. This was not the case at all – every single paraphrase used “want” or a closely related term, like “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 the supervenience relation it bears 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. Clearly such states belong to Boeing rather than the individuals whose actions such non-mental states nonetheless supervene on. Would people paraphrase group non-mental state ascriptions and group mental state ascriptions in the same way, and differently from the way in which they paraphrase ascriptions of mental states to individuals? Or, would people paraphrase group mental state ascriptions differently from the other two kinds of ascriptions, and indicatively of a non-realist approach to such ascriptions? Or, as Knobe and Prinz would seem to predict, would there

be a parallel between people’s paraphrases of intentional state and non-mental state attributions to groups and persons, but no such parallel for group phenomenal state attributions? 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.

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”.¹² For each participant we calculated an “anti-realist score” by awarding 1 point each time the participant selected “they” to replace the underlined name in one of our test sentences, and 0 points each time the participant

¹² Versions of our test materials are included in the Appendix.

selected “it”. Given that we had four test sentences, a score of 4 is indicative of extreme anti-realist sentiments, while a score of 0 indicates realism about the particular kind of group state ascriptions.¹³

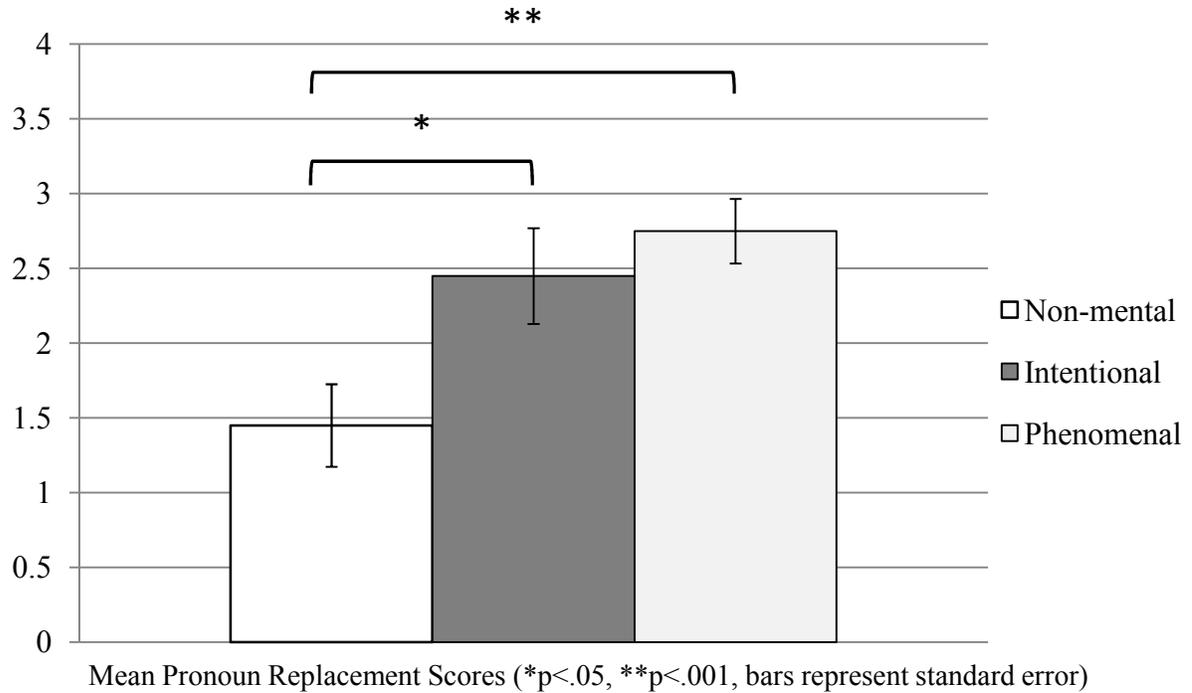
Our prediction was that participants would be more inclined to use the “it” pronoun in the condition involving non-mental state ascriptions than in either mental state condition. That is, we suspected that people were distributivists about both kinds of group mental state sentences (i.e., understanding 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. Participants were significantly more likely to select the plural pronoun paraphrase for intentional mental state ascriptions (mean anti-realist score: 2.45) than they were for non-mental state ascriptions ($m = 1.45$).¹⁴ 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.¹⁵ However, no significant difference emerged between plural pronoun selections for intentional and phenomenal state ascriptions.¹⁶ Sentences were counter-balanced, and no order effects emerged.

¹³ Actually, some of those who chose “it” for each group name may simply have been committed to strict grammatical rules. Since it is impossible to tell what the true motivation was for these purists, and since, in any case, their inclusion could only hurt our cause, we did not exclude them from our analyses.

¹⁴ 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

¹⁵ $t(40) = 3.7$, $p < 0.001$ (two-tailed), SD (phenomenal) 0.967, *Cohen’s d* = 1.17

¹⁶ $t(40) = 0.777$ $p = 0.422$ (two-tailed), *Cohen’s d* = 0.24



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

Non-mental: 1

Intentional: 3

Phenomenal: 2.5¹⁷

People most frequently chose “they” to paraphrase a group name when the sentence involved an ascription of *either* an intentional *or* a phenomenal mental state to the group. They most frequently chose “it” when the sentence involved a non-mental state ascription. This suggests that the way people understand both varieties of group mental state ascriptions—as revealed by the pronouns they use to paraphrase these sentences—is quite different from the way they understand ascriptions of mental states to persons. After all, even though it’s popularly

¹⁷ The median score here is 2.5 because there was an even number of subjects, and the two scores in the middle of the range were 2 and 3.

recognized that a person’s mental states supervene on their neural states, we never paraphrase ascriptions of mental states to individuals using plural pronouns.

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. Our participants are not making assessments about what sounds natural—assessments that are likely to be pushed around by an understanding of the conversational pressures that could motivate a given response. Nor are participants’ responses subject to concerns of brevity. Participants are choosing between the 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: Explicit Paraphrase

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,” and “How does everyone like their soup?” ‘They’ can also be used to refer to a collective (rather than to the members that constitute it), as when one

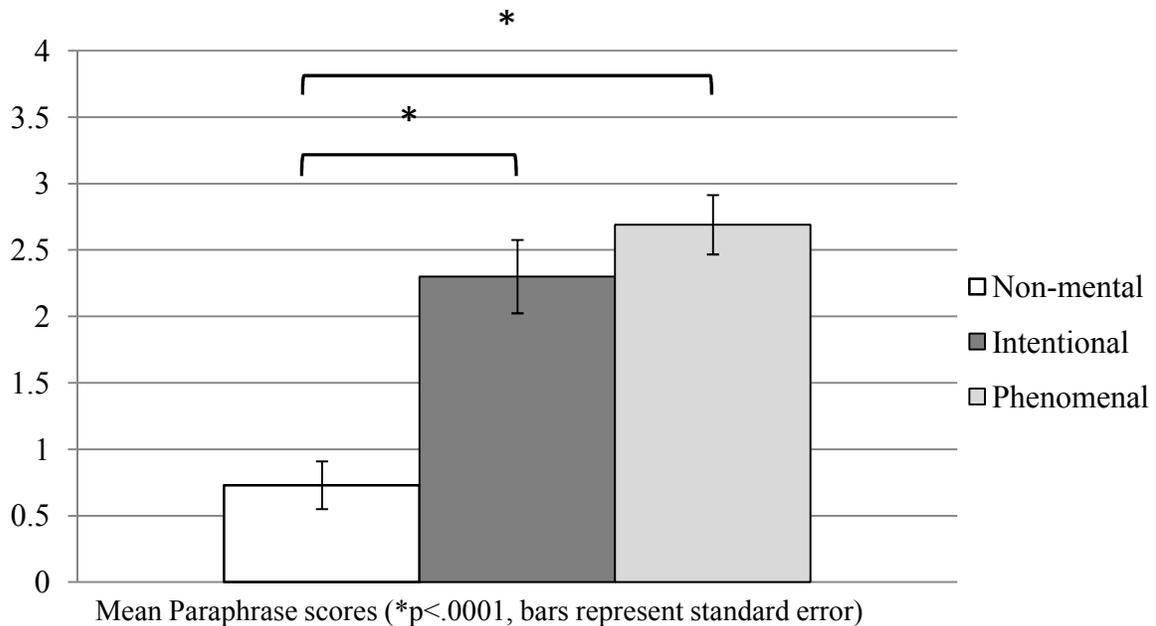
says of a truck-load of shampoo bottles, “They weigh a ton.” Perhaps ‘they’ is being used in the context of our study in some such singular-substitutive way. Of course, such a challenge to our findings would not by itself explain the asymmetry between mental and non-mental ascriptions that we found in 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 attributions.

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 (see Appendix). 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 ascription. One of these sentences was *distributivist*, in that it described the ascription as being directed at some relevant subset of members in the group; and one of the sentences 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*. Scores were tabulated by assigning zero points for each *realist* selection and one point for each *distributivist*

selection, then summing the points for all four sentences. Thus, a strong realist score would be 0, while a strong distributivist score would be 4.

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 a significant differences both between phenomenal and non-mental conditions¹⁸ and between intentional and non-mental conditions.¹⁹ However, responses to phenomenal and intentional conditions were not significantly different.²⁰



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

Non-mental: 0

¹⁸ $t(53) = 6.743, p < .0001, SD(\text{phenomenal}) 1.198, SD(\text{non-mental}) .92, \text{Cohen's } d = 1.85$

¹⁹ $t(51) = 4.706, p < .0001, SD(\text{intentional}) 1.44, \text{Cohen's } d = 1.32$

²⁰ $t(54) = 1.116, p = .269, \text{Cohen's } d = 0.30$

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, i.e., as a kind of shorthand for attributions to the members (or some relevant subset of members) that compose the group. 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.

4. Who We Think Thinks

Our discussion thus far has implications both methodological and specific to the topic of mental state attribution. Let us first emphasize the specific issue of critical discontinuity. Is there actually a critical discontinuity between the properties that dispose people to attribute phenomenally conscious states and intentional states? 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 verbal ascriptions of group mental states has been essential to the case for the discontinuity thesis. Thus, our studies undercut important evidence in favor of a critical discontinuity. If 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.

Of course, even conceding that asymmetries in appropriateness ratings for sentences do not, in light of our evidence, support a critical discontinuity in folk metaphysics of mind, one may wonder how we explain the apparent asymmetry in appropriateness 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.’ How, then, can we accommodate the intuitiveness of Knobe and Prinz’s results while maintaining the distributivist reading of folk interpretations?

On the one hand, we suspect that evidence for this asymmetry itself may not be as strong as some have suggested. Knobe and Prinz’s study materials in particular have garnered specific criticisms. Arico (2010), for example, points out that Knobe and Prinz’s prompts are not minimal pairs and presents empirical evidence that the asymmetry disappears once you introduce equally contextualized intentional and phenomenal prompts of equal length. Still, it is hard to resist the intuition that, “Microsoft believes certain things,” sounds more natural (in some sense of the term) than, “Microsoft feels sad about certain things.” How can we explain this, assuming, as we have argued, that both utterances are understood to encode attributions of states to human members, equally able both to believe and feel sad?

We suggest—though we will not explore the suggestion in any detail here—that such an intuitive asymmetry is (if it is real) the effect of a tendency to regard verbal ascriptions of states to groups appropriate insofar as the state is one that members of the group occupy *qua members* of the group. Since a corporate employee is often thought to exercise intentional states in the execution of his job, intentional state ascriptions will tend to be thought appropriate, according to this *qua member’s* (QM) principle. It is not in the job description of a Microsoft employee to feel sad, so ascriptions of sadness to Microsoft, though really only shorthand attributions of a

phenomenal state to human employees, will be thought inappropriate. If the QM-principle is correct then there could, in principle, be groups for whom phenomenal state ascriptions would be deemed appropriate. (Perhaps Al Qaeda is a plausible candidate.) In any case, this tentative suggestion requires serious empirical investigation before it can be defended with any certainty.²¹

However one explains the purported asymmetry in appropriateness assessments, our results undermine evidence offered in favor of the discontinuity view. But the results fit perfectly with what would be predicted by the continuity view. Notice that, in both 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 treat the two kinds of group mental state attributions the same. But this result conforms best with a continuity view, according to which the same sorts of properties are relevant for determining whether an entity has phenomenal or intentional states.

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 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. But then Knobe and Prinz's finding that experimental participants assent to intentional, but not

²¹ An alternative, but related, proposal emphasizes the contextual situation of the attribution. A group is generally situated within a relatively specific context in virtue of the type of group it is (corporations within business contexts, athletic teams within sports contexts, etc.). Attributions to groups may be judged appropriate in part due to the fittingness of the attribution within the relevant context. Following the illustration above, sadness is not typically appropriate in a professional/business context, although there are specific professional/business contexts in which feeling sad may be appropriate (e.g., stock market collapses or extensive layoffs).

phenomenal, verbal ascriptions to groups cannot support a discontinuity related to biological body, for, after all, *group members* are the self-same embodied *people* who are thought to have phenomenal states if any entities are. 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 into group mental states, but also of research focusing on attributions of mental states to non-biological entities other than groups, as two prominent examples demonstrate. 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). In a different study, experimental philosophers Justin Sytsma and Edouard Machery (2010) describe a situation in which a simple robot follows instructions to move a red box (rather than a green or blue one). They then asked participants whether the robot saw red. Participants recorded their responses on a 7-point scale anchored at 1 with “clearly no,” at 4 with “not sure,” and at 7 with “clearly yes.” Lay-persons tended more towards responding “clearly yes,” with a mean response of 5.15, whereas philosophers had a mean score of 3.48. Sytsma and Machery conclude that these “results indicate that philosophers’ concept of phenomenal consciousness is *not* how the folk understand subjective experience” (308).

In these studies researchers assume that participants’ verbal responses to questionnaires are indicative of fairly specific beliefs. Thus these researchers exercise some degree of interpretive authority on how participants understand various kinds of sentences. However, the problems we have identified in Knobe and Prinz’s argument (2008) obviously raise the possibility of similar problems for these other studies as well. 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.”²² Sytsma and Machery’s findings that people attribute color vision to robots, on the other hand, may reflect the pressure to communicate the fact that the robot is obviously processing color-related information, rather than a judgment that robots and humans have a common capacity for seeing color. The philosophers’ concern with literalness and verbal precision may explain their comparative reluctance to go in for this strictly inaccurate gloss.²³

It is natural to enquire as to the status of the kind of inaccurate talk that we have been exploring. In particular, when people say, “Boeing felt anxious about having to lay off workers,” but mean, *some relevantly important or large subset of Boeing employees felt anxious about the fact that HR was going to have to lay off some workers*, is it correct to say that people speak figuratively? As we will explain below, our answer is, yes and no, and, as we will argue, being careful to distinguish different notions of figurative language can have important implications for cognitive scientific research.

²² Bloom and Veres transcribe only one forward group video description and one reverse description, so we cannot observe any general trend. However, it is interesting to note that the forward description, glossed in intentional terms, is only 85 words long, whereas the reverse description, which per necessity relies more on trajectory talk, is, at 147 words, almost twice as long.

²³ Sytsma and Machery consider a distinct, though related, response to their finding, that the expression “seeing red” is ambiguous between an informational and a phenomenal reading, and that the folk tend to give it the informational reading whereas philosophers give it the phenomenal one. Our point here is instead that, wanting to express what they think (that the robot is processing color-related information) but being unable to given the restricted survey choices, it would be reasonable for participants to opt for a response that, while perhaps not completely accurate, gets the point across more or less. Philosophers on the other hand, being trained for verbal precision, are less inclined to such loose talk.

In a previous study, some of the authors of the current paper (Arico et al., 2011) set out to show that subjects regarded attributions of mental states to groups as figurative, rather than literal, claims. 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. Contrary to the authors’ prediction, participants claimed that attributions of such intentional mental states to groups were *literally* true. How can this be if, as we have claimed above, such uses are ones in which speakers actually mean to attribute the mental states to the employees of Boeing, not to the corporation itself, over and above its members?

The crux of our response to the previous findings is this: There exist two notions of figurative language. There is an ordinary language notion of the figurative, often explicitly flagged as “metaphorical”. For example, in news reports, the campaign leading up to the conference of the African National Congress is described as, “metaphorically speaking, dirty and very bloody.” Whereas the co-founder of Sun Microsystems is characterized as having, “pitched a tent, metaphorically speaking,...and preached a new testament vision of golf.”²⁴ These utterances are imagistic, visceral, or otherwise striking. But there also exists a notion of the figurative common to philosophers and linguists which encompasses these utterances and more. The philosopher’s concept of the figurative encompasses all utterances involving loose use of terms.

²⁴ The examples come, respectively, from *Interpress Service News*, “Politics South Africa” (12/18/2007) and from the *Wall Street Journal*, “The Radical New Vision of Golf,” (2/11/2011).

Whatever the ordinary notion of figurative language involves, the results from Arico 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 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.²⁵ 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.²⁶ 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.²⁷ In so far as we are drawing conclusions about what people *actually* think from what people say, it will be important

²⁵ As in section 2 (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. Perhaps “Acme Corp” simply refers to the employees and share-holders, or, more likely, it is semantically context sensitive. In either case, we cannot assume that people really think that group entities can have intentional mental states *even if* we know people are not departing from encoded word meaning.

²⁶ 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.

²⁷ In other work, one of the authors of this paper (Phelan) has argued that it involves a certain potential to affect a hearer in a distinctive way. See Phelan, MS, for further discussion.

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.²⁸

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 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.

²⁸ 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. As we pointed out above, various theories of communication recognize that conversational pressures can lead us to pronounce a sentence whose strict meaning (the strict meaning of the pronounced elements, that is) differs from what we mean in speaking. Depending on which theory of communication one embraces, one will appeal to different mechanisms, such as loose use or implicature, in explaining how, in instances such as these, speakers convey the propositions they actually and strictly endorse.

Appendix: Test Sentences and Paraphrases

Study 1: Pronoun Replacement Task

In each condition of this study, there were four test sentences and four filler sentences. Here we reproduce the test sentences for each condition. Complete versions of test materials are available online at: <http://epl.web.arizona.edu/BoeingStats.html>

Non-Mental:

1. After Phi Lambda Mu received the announcement that the ban on large house parties had been lifted, Phi Lambda Mu had a larger than expected pledge class.
2. When MADD's Drunk Driving Prevention Act failed, MADD lost donations.
3. After Boeing lost the Army contract, Boeing needed to lay off workers.
4. After the existence of the Secret Task Force was revealed, the Secret Task Force met with public condemnation.

Intentional:

1. After Phi Lambda Mu received the announcement that the ban on large house parties had been lifted, Phi Lambda Mu assumed it was okay to have keg parties again.
2. When MADD's Drunk Driving Prevention Act failed, MADD contemplated alternatives.
3. After Boeing lost the Army contract, Boeing expected to have to lay off workers.
4. After the existence of the Secret Task Force was revealed, the Secret Task Force wanted to minimize public reaction.

Phenomenal:

1. After Phi Lambda Mu received the announcement that the ban on large house parties had been lifted, Phi Lambda Mu relished the prospect of having keg parties again.
2. When MADD's Drunk Driving Prevention Act failed, MADD got extremely depressed.
3. After Boeing lost the Army contract, Boeing felt anxious about having to lay off workers.
4. After the existence of the Secret Task Force was revealed, the Secret Task Force felt threatened by the public reaction.

Study 2: Forced Choice Paraphrase Task

Non-Mental:

After Boeing Company lost the Army contract, **Boeing Company needed new clients.**

Which of the following do you think best **describes** what is meant in the **BOLD portion of the above sentence:**

_____ The relevant employees at Boeing Company needed new clients.

_____ Boeing Company itself, not just its employees, needed new clients.

When MADD's Drunk Driving Prevention Act failed, **MADD lost donations.**

Which of the following do you think best **describes** what is meant in the **BOLD portion of the above sentence:**

_____ The majority of the organization's relevant members lost donations.

_____ The organization itself, not just its members, lost donations.

After Phi Lambda Mu received the announcement that the ban on large house parties had been lifted, **Phi Lambda Mu had a larger than expected pledge class.**

Which of the following do you think best **describes** what is meant in the **BOLD portion of the above sentence:**

_____ The fraternity itself, not just the brothers, had a larger than expected pledge class.

_____ The majority of relevant fraternity brothers had a larger than expected pledge class.

After the existence of the Secret Task Force was revealed, **the Secret Task Force encountered public resistance.**

Which of the following do you think best **describes** what is meant in the **BOLD portion of the above sentence:**

_____ The majority of relevant task force participants encountered public resistance.

_____ The task force itself, not just its participants, encountered public resistance.

Intentional:

After Boeing Co. lost the Army contract, **Boeing Co. expected layoffs would be necessary.**

Which of the following do you think best **describes** what is meant in the **BOLD portion of the above sentence**:

_____ Boeing Company itself, not just the employees, expected that layoffs would be necessary.

_____ The majority of the relevant employees at Boeing Company expected that layoffs would be necessary.

When MADD's Drunk Driving Prevention Act failed, **MADD contemplated alternatives.**

Which of the following do you think best **describes** what is meant in the **BOLD portion of the above sentence**:

_____ The majority of the relevant members of the organization thought about alternative ways of getting similar legislation passed.

_____ The organization itself, not just the members, thought about alternative ways of getting similar legislation passed.

After Phi Lambda Mu received the announcement that the ban on large house parties had been lifted, **Phi Lambda Mu assumed it was okay to have keg parties again.**

Which of the following do you think best **describes** what is meant in the **BOLD portion of the above sentence**:

_____ The fraternity itself, not just its members, contemplated the announcement and determined that it was okay to have keg parties again.

_____ A majority of the relevant fraternity brothers contemplated the announcement and determined it was okay to have keg parties again.

After the existence of the Secret Task Force was revealed, **the Secret Task Force wanted to minimize the public reaction.**

Which of the following do you think best **describes** what is meant in the **BOLD portion of the above sentence**:

_____ The majority of the relevant task force participants wanted to minimize public reaction.

_____ The task force itself, not just the participants, wanted to minimize public reaction.

Phenomenal:

After Boeing Co. lost the Army contract, **Boeing Co. felt anxious about layoffs being necessary.**

Which of the following do you think best **describes** what is meant in the **BOLD portion of the above sentence:**

_____ Boeing Company itself, not just the employees, experienced some mild nausea and restlessness about potential layoffs.

_____ The majority of the relevant employees at Boeing Company experienced some mild nausea and restlessness about potential layoffs.

When MADD's Drunk Driving Prevention Act failed, **MADD got extremely depressed.**

Which of the following do you think best **describes** what is meant in the **BOLD portion of the above sentence:**

_____ The majority of the relevant members of the organization were overcome with sadness and feelings of helpless despair.

_____ The organization itself, not just the members, was overcome with sadness and feelings of helpless despair.

After Phi Lambda Mu received the announcement that the ban on large house parties had been lifted, **Phi Lambda Mu relished the prospect of having keg parties again.**

Which of the following do you think best **describes** what is meant in the **BOLD portion of the above sentence:**

_____ The fraternity itself, not just the brothers, felt a wave of excited anticipation at the thought of having keg parties again.

_____ The majority of the relevant fraternity brothers felt a wave of excited anticipation at the thought of having keg parties again.

After the existence of the Secret Task Force was revealed, **the Secret Task Force felt threatened by the public reaction.**

Which of the following do you think best **describes** what is meant in the **BOLD portion of the above sentence:**

_____ The majority of relevant task force participants felt threatened by public reaction.

_____ The task force itself, not just the participants, felt threatened by public reaction.

Works Cited:

- Arico, A. unpublished manuscript. “Deregulating Corporate Consciousness: A Critique of Knobe and Prinz’s *Intuitions About Consciousness*.”
- 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. 1994. “Conversational Implicature.” *Mind and Language* 9: 124-62.
- Bach, K. 2001. “You Don’t Say?” *Synthese* 128: 15--44.
- Block, N. 1978. “Troubles with Functionalism.” *Minnesota Studies in the Philosophy of Science*, Volume IX. C. Wade Savage ed. Minneapolis: University of Minneapolis Press, 261-325.
- 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.
- Brentano, F. 1973. *Psychology from an Empirical Standpoint*. New York: Humanities Press.
- Cullen, S. 2009. “Survey-Driven Romanticism.” *Review of Philosophy and Psychology* 1: 275-296.
- Fiala, B., Arico, A., and Nichols, S. forthcoming. “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.
- Gilbert, M. 1996. *On Social Facts*. Princeton, NJ: Princeton University Press.
- Gopnik, A. and Astington, A. W. 1988. “Children’s understanding of representational change and its relation to the understanding of false belief and the appearance-reality distinction.” *Child Development*, 59, 26-37.
- 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.
- Grice, P. 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Heider, F. and Simmel M. 1944. "An experimental study of apparent behavior." *American Journal of Psychology* 57, 243–259.
- 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.
- Knobe, J. 2008. "Experimental Philosophy of Consciousness." *Scientific American: Mind*
- 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.
- Pettit, P. 1993. *The common mind: An essay on psychology, society and politics*. New York: Oxford University Press.
- . 2003. "Groups with minds of their own." In *Socializing Metaphysics*, edited by F. Schmitt. New York: Rowman & Littlefield
- Phelan, M. 2010. "The inadequacy of paraphrase is the dogma of metaphor." *Pacific Philosophical Quarterly* 91: 481-506.
- Phelan, M. MS. "Making the metaphor move: Differentiating figurative and literal language."
- 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.

Sperber, D. and Wilson, D. 1995. *Relevance: Communication and cognition, 2nd Edition*. Oxford: Blackwell.

Stanley, J. and Szabó, Z. 2000. "On Quantifier Domain Restriction," *Mind and Language*, 15/2: 219-261.

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