

Phenomenal Consciousness Disembodied*

WESLEY BUCKWALTER
University of Waterloo
wesleybuckwalter@gmail.com

MARK PHELAN
Lawrence University
mark.phelan@lawrence.edu

Abstract: We evaluate the role of embodiment in ordinary mental state ascriptions. Presented are five experiments on phenomenal state ascriptions to disembodied entities such as ghosts and spirits. Results suggest that biological embodiment is not a central principle of folk psychology guiding ascriptions of phenomenal consciousness. By contrast, results continue to support the important role of functional considerations in theory of mind judgments.

1. *Introducing Some Phenomenal Bodies*

“Suppose we convert the government of China to functionalism, and we convince its officials to realize a human mind for an hour.” This is, of course, the beginning of Ned Block's famed “Nation of

* This is the penultimate version of the paper to appear in *Continuum Advances in Experimental Philosophy of Mind*. Please cite the final, published version if possible. The authors are grateful to Matthew Kieran, Joshua Knobe, Aaron Meskin, Shaun Nichols, Justin Sytsma, John Turri, and Joshua Weisberg for helpful comments on earlier versions of this paper.

China” thought experiment (1978, p. 279). In it, Block asks us to imagine that approximately 1 billion people (roughly the population of China at the time) come together to simulate the inner workings of a normal human brain. Each person is given a two-way radio that enables him or her to communicate with others, much like individual neurons in the brain communicate with other neurons. What's more, this complex communication network is hooked up to a remote, artificial body. As the sensory organs of that body are stimulated, the external state of the body is reported on a system of satellites, visible anywhere within China.

According to a set of specified rules, the Chinese respond to these satellite reports by relaying information and commands to one another, and issuing instructions to the body that cause it to execute various behavioral routines. In a sense, the nation of China has now become a “China-Brain” hooked up to an artificial body (just by radio waves rather than electrochemical impulses). Arguably, such a system satisfies a purely functional description of a mind. Broadly speaking, functionalism defines individual mental states in terms of sensory inputs, behavioral outputs, and relations to other mental states. So according to this view, the various states of the nation of China would amount to the thoughts, feelings, desires, and so on, of normal human beings.

But could a being of such odd construction be said to have a mind, filled with the same kinds of thoughts, feelings, and desires as us? Part of the original goal of Block's “Nation of China” thought

experiment is to demonstrate that it could not. In fact, Block finds it doubtful that we will say that the “China-Brain” has any mental states at all—least of all any “qualitative states, raw feels, or immediate phenomenological qualities” (p. 281). It simply is not made of the right stuff for such states to be possible. This intuition has been taken by some as evidence against the claim that certain theories of functionalism adequately capture how the mind works.

The significance of Block's thought experiment has been debated. But the basic example has remained influential in both psychology and cognitive science.¹ One prominent position that has emerged from this discussion is the view that only certain sorts of entities are capable of having phenomenally conscious mental states such as emotions or raw feels. In addition to being functionally organized in the right sort of way, entities capable of certain mental states (in particular, phenomenal experiences) it is said, must also have the *right sort of biological bodies*. They must be made of the right stuff.

This traditional view held by certain philosophers of mind about how the mind works has also been explored by others as an important principle of folk psychology guiding ordinary mental state ascription. A number of researchers in philosophy, psychology, and cognitive science have defended weaker or stronger ver-

¹ Much the same could be said for both the original purpose and subsequent legacy of Searle's “Chinese Room” argument (1980) with respect to the questions of function and embodiment.

sions of a view that we will call the *embodiment hypothesis* (Knobe 2008, Knobe & Prinz 2008, Gray et al. 2011). Generally speaking the embodiment hypothesis states that unified *biological embodiment* is a major psychological factor that cues ordinary attribution of experiences, feelings, emotions, and so on, to other entities. The strongest version of this view is that phenomenal attribution requires biological embodiment. Weaker versions focus on relative levels of attribution, claiming that phenomenal attributions are more likely to be cued as an entity's biological body becomes more salient.²

To begin to get a sense of why many have been drawn to the embodiment hypothesis, it may be helpful to consider a specific example of some work on the role of embodiment on mental state ascriptions to different sorts of entities. For instance, in their work studying people's intuitions about group agents, Knobe & Prinz (2008) observe:

It is a striking fact about group agents that we ascribe to them some types of mental states but not others. We might

² We use the terms 'phenomenal consciousness' and 'phenomenal attribution' throughout the paper when referring to states typically classified by philosophers as qualitative states or states of subjective experience. However there are some doubts in the experimental philosophy of mind literature concerning whether people have the concept of phenomenal consciousness (see Sytsma & Machery 2010). We emphasize that none of our main arguments or findings here depend on whether or not non-philosophers draw the phenomenal/non-phenomenal distinction when ascribing experiences, feelings or emotions to disembodied entities, and set the issue aside. We thank Justin Sytsma for discussion on this point.

say that Microsoft intends something or wants something or believes something...but there are other kinds of ascriptions that we would never make to Microsoft. For example, we would never say that Microsoft was feeling depressed. (p. 73)

Knobe & Prinz conduct several studies and find that this is indeed the case. People are very reluctant to ascribe states like *feeling depressed* to the Microsoft Corporation. They go on to explain this striking fact by appealing to two claims. The first claim is that there are important differences in how people ascribe intentional states (like intending or wanting) on the one hand, and states requiring phenomenal consciousness (like feeling sad or depressed) on the other. The second claim is that attributions of these latter kinds of mental states are “sensitive in a special way to information about physical constitution” (p. 73).³

Strictly speaking, the Microsoft Corporation does have a *physical* body. It has a body in the sense that it has a physical presence. It is comprised of factories built of brick and mortar, office buildings, technical laboratories, as well as researchers and employees spread out all across the globe. But the Microsoft Corporation obviously doesn't have a unified body. It is spatially disconnected and includes many disparate kinds of parts. And while it has individual members that are human, the Microsoft Corporation itself

³ For alternative explanations of Knobe & Prinz's findings, see Phelan et al. (forthcoming), Sytsma & Machery (2009), and our discussion in Section 3.5 below.

clearly lacks a biological body. So while it has a body in some extended sense, it lacks a unified body comprised, among other things, of flesh and blood. According to Knobe and Prinz, we are reluctant to attribute phenomenal states or subjective experiences to the Microsoft Corporation because Microsoft lacks the right kind of unified biological body.⁴

As far as we know, the embodiment hypothesis about folk psychological judgments has not been endorsed by Block or other traditional philosophers of mind directly. Knobe & Prinz (2008) argue for the hypothesis insofar as facts about physical constitution can explain low phenomenal state ascriptions to group agents. Knobe (2008) and Gray et al. (2011) argue for the hypothesis on the grounds that body salience correlates with higher attributions of phenomenal capacities.⁵ But leaving aside questions of exactly who has endorsed specific versions of the embodiment view, we note that work by Block, Knobe, and others has made the general embodiment hypothesis a very attractive view that theorists in both psychology and philosophy of mind might be tempted to accept. The general view that unified biological embodiment is a major psychological factor that cues ordinary attribution of experiences, feel-

⁴ Knobe and Prinz seem to focus on the disunity of *corporate* entities as the crucial factor, since one of their later studies suggests that an enchanted chair with a unified “body” can have phenomenal states.

⁵ This work focuses on the psychological cues for the attribution of mental capacities, while the experiments we present below focus on the cues for attribution of specific mental states. More research is needed to study the subtle differences between these two closely related research questions.

ings, emotions, and so on seems to fit some of the data that has been collected across a number of influential studies on mental state ascription. The view would predict that entities with the right kind of biological body are the ones typically thought capable of having phenomenally conscious experiential states. And conversely it holds that entities without the right kind of biological body, such as robots, groups, and ghosts, are typically attributed phenomenal mental states at only very low levels. Furthermore, the embodiment view would be one way of *explaining* the psychological basis of intuitions such as those in Block's original thought experiment, which some philosophers have used as evidence for key philosophical conclusions in the metaphysics of mind.⁶

Nonetheless, we think that philosophers and psychologists should be slow to accept the general embodiment view. Some of the key thought experiments and empirical studies that have been presented to date suggest that embodiment actually does not play that important of a role in the way we ordinarily attribute mental states. Specifically, a number of philosophers and cognitive scientists sympathetic to *functional accounts* of the mind have suggested that intuitions favouring embodiment in both thought experiments, like Block's China Brain, as well as empirical studies on mental state ascription might in fact be trading on subtle cues and distract-

⁶ After all, the nation of China is one (special kind) of group entity. (Though see Phelan et al., forthcoming, for an independent source of resistance to thinking that groups really have minds.)

tions related to the *functional organization* of the target entities. On the philosophical side, Dennett (1991), for example, argues that “China-Brain” does not constitute an acceptable counterexample to functionalism. Instead, he argues that Block’s thought experiment unfairly relies on a “misdirection of the imagination” because it nonchalantly invites readers to buy into the unlikely idea that the China-brain is complex enough to satisfy the functional roles associated with particular mental states. And similarly, on the empirical side, we have argued elsewhere (Phelan and Buckwalter, forthcoming), that many of the experimental materials researchers have used to study the influence of embodiment on mental state ascriptions include potential confounds. For example, many include subtle but crucial functional descriptions of entities and their environments (e.g., information about inputs, outputs, and other mental states to which the entities are subject). In addition to considerations about unified biological embodiment, the inclusion of these potential confounds makes it difficult to assess existing research purported to support the embodiment hypothesis.

Before continuing, it may be helpful to pause and consider how functional information might influence mental state ascriptions. Suppose, for instance, someone was trying to figure out whether or not an entity (let’s call this entity ‘Bob’) feels happiness or anger about some state of affairs (such as current low interest rates). In assessing whether Bob is happy or angry about low interest rates, information about Bob’s other mental states will be im-

portant. Bob is more likely to be happy if he wants to borrow a large sum of money; more likely to be angry if he wants to make money as a lender. Our assessment of Bob's emotional state will also be affected by our beliefs concerning the external stimuli to which Bob is subject. If for instance Bob is angry about the low rates, but then hears a newscast that they've just increased, we will likely temper our assessment of Bob's anger accordingly. Finally, Bob's overt behavior will factor into our assessment of his mental states. If we see Bob cursing or tearing up his lender's agreement, we will be more likely to conclude Bob is angry over the low interest rates. It strikes us as obvious that such functional information contributes to mental state assessment. Indeed, we have shown in prior work that people often consult these kinds of cues when deciding whether or not to ascribe phenomenal states (Buckwalter and Phelan, 2012).

Of course, even if functional information of this sort does cue phenomenal state attributions, it could still be that embodiment (for instance, whether 'Bob' is a normal human being, a group, or an immaterial ghost) constitutes an important ascription cue as well. However, we think there are good reasons to reject the embodiment hypothesis. This is what we will attempt to demonstrate. In the remainder of this paper, we examine people's ascriptions of experiential states to entities lacking a biological body. Our goal is to see if ascriptions of phenomenal states to these sorts of entities differ from ascriptions made about normal human beings, or if they tend to work in the same basic way.

For this task, a number of different disembodied entities might have been used. We chose to begin our examination of disembodied ascription with the *phantasmally* disembodied—ghosts and spirits. Lacking in any body whatsoever, spirits constitute the ultimate test of the embodiment view. If the embodiment hypothesis were correct, and embodiment were a crucial cue for phenomenal state attribution, then we would expect important differences in ascription between human beings, on the one hand, and disembodied ghosts and spirits, on the other—just as we expect to find important differences in phenomenal state attribution for functional information. If functional information—information about the goals, desires, etc, of an entity—tends to cue mental state ascription independently of whether the entity has a unified biological body, then it undermines the embodiment hypothesis. This is what we set out to investigate, using spirits as our medium.

2. *Previous Studies on Mental State Attributions to Disembodied Entities*

We are under no illusion that our investigation into how people attribute mental states to the disembodied is unprecedented. Several influential studies on God and ghosts have already been conducted. Indeed, previous findings seem to offer contradictory evaluations of the embodiment hypothesis. Some have taken the findings of Gray,

Gray and Wegner (2007) to support embodiment.⁷ Gray et al. analyzed comparative attributions of a range of mental states to a cast of “characters” ranging from babies to adults, from robots to animals to the dead, and including the ultimate disembodied entity: God. They found that people were less willing to attribute phenomenal mental states (such as feeling fear, hunger, or pain) to God than to many of the other characters in their study. This finding is consistent with the embodiment view, since, presumably, God is thought to lack a body whereas other of the characters are not. Furthermore, insofar as lack of embodiment is what explains these low attributions to God, the findings are inconsistent with a simple functionalist account of people's phenomenal state attributions. However, as we have argued elsewhere (Phelan and Buckwalter, forthcoming), the findings are readily explicable in functionalist terms, since God is thought to be the ultimate being, who wants for nothing. Surely then he will be thought to suffer fear, hunger, and pain less often than a child or a toad (two other of Gray et al.'s characters).⁸

On the other hand, findings from Jesse Bering and colleagues could be interpreted as challenging the embodiment view. Bering found that adults thought psychological functions—including emotional states—continued after biological death in an agent killed on his daily commute (2002). Bering and Bjorklund (2004)

⁷ Note however that Gray & Wegner (2010) question whether these prior findings about God are best interpreted as supporting embodiment.

⁸ Similar considerations, we think, explain Gray et al.'s findings for other phenomenal states.

found a similar pattern for children, who continued to attribute emotional and other mental states to a mouse after it was eaten by an alligator. And Bering et al. (2005) found that both secularly and parochially educated children under 10 were proportionally more likely to disagree with statements indicating that psychological functions including emotional states ceased at death. On the assumption that each population thought of the recently deceased as disembodied (a supposition supported by the fact that each population tended to think that biological function ceased at death), these studies present *prima facie* counter-evidence to the embodiment view.

While Bering's work is illustrative, for our present purposes it does not constitute a true test of the embodiment hypothesis. For one thing, it doesn't explicitly compare people's attributions of phenomenal states to the disembodied with their attributions of phenomenal states to normal, embodied humans. Thus it might have missed a tendency to attribute mental states in a way consonant with embodiment. For another thing, it doesn't explicitly manipulate function. Thus it offers no comparison between cues related to embodiment and other salient cues of phenomenal state attribution. Finally, Bering's experimental materials don't explicitly inform participants that the recently deceased agents are disembodied, nor do they ask participants whether they conceive of the dead agents in this way. It thus remains a possibility that experimental participants are not equating death with disembodiment in a way that

would shed light on the embodiment hypothesis. Therefore we use this body of research on ordinary beliefs about souls as a point of departure for testing the embodiment view.⁹

3. *Disembodying Ascription*

We present five experiments investigating people's willingness to ascribe emotional states to disembodied ghosts and spirits. The take home message of this section is that there are certain conditions under which people are perfectly willing to ascribe phenomenal states to these kinds of disembodied, non-biological entities. Using functional information as a comparison, Experiment 1 demonstrates that ascriptions of emotion to ghosts matches that of ordinary human beings, so long as either of those entities satisfy the relevant functional descriptions. Experiment 2 replicates this effect using experimental probes that lack intentional objects with potentially biasing contextual information (see Section 2.2). Experiments 3-4 display commensurate results for attributions to eternally disembodied spirits. And lastly, using an explicit comparison technique, Experiment 5 provides evidence that in making such ascriptions,

⁹ These are not criticisms of Bering's work, since he wasn't out to investigate the issue of physical realizers at all. In fact, Bering is one of a number of theorists arguing for a particular view about the source of afterlife beliefs, which Bering (2011) calls the *simulation constraint hypothesis*. Nichols (2007), another proponent of the view, encapsulates the basic idea as follows: "part of the reason we believe in immortality is that we can't imagine our own nonexistence" (p. 216). Interesting as the connections are between the embodiment hypothesis and afterlife belief, we set them aside.

participants literally attribute the same phenomenal states to disembodied ghosts and spirits as they attribute to ordinary human beings.

3.1. Experiment 1: Disembodied Ghost

3.1.1. Methods

We begin with a between-subjects multifactor experiment designed to test the influence of embodiment and functional cues on ascriptions of phenomenal states.¹⁰ Participants in Experiment 1 (N=158, 85 female, median age = 27) were presented with the following story about Bob, his ex-wife Melissa, and their son Henry. Roughly half of the participants received a version where the functional information (in this case, the goal Bob attempts to bring about) is to cause his son Henry to hate his mom, while the other half saw a condition in which Bob aimed to make Henry happy:

[ANGER/HAPPY] Bob and Melissa have been married for 15 years. After several months of intense bickering and fighting, they decide to get a divorce. Bob moves out of the house, but still tries to spend time with Henry, their ten-year-old son. He also continues to keep close tabs on his ex-wife Melissa. One day, Bob learns that Melissa has started a new romantic relationship.

¹⁰ Participants in this and all subsequent experiments were recruited and tested using commercially available online platforms (Qualtrics and Amazon Mechanical Turk). Participants were located in the United States, and over 85% reported English as a first language. They were paid between \$0.30 and \$0.45 for their participation. Participants were prohibited from taking more than one study.

He hires a private investigator to follow the couple, and take photos of them over a romantic dinner. Bob knows that it will [*cause Henry to hate his mom, Melissa, if he learns that she/ make Henry incredibly happy if he learns that his mom, Melissa,*] has started a new, meaningful relationship. Suddenly, Bob gets an idea. If he leaves the pictures in Henry's treehouse in the backyard, Henry is sure to find them when he gets home from school that day. So, Bob jumps in his car and drives to Melissa's house.

After reading one of the versions above, half of the participants saw a conclusion to the story where Bob's biological body is made salient:

[EMBODIED] On the drive over however, Bob is in a car accident. Bob emerges from his car and looks over his body. Everything seems to be completely fine—his head, legs and arms. But even though Bob has been in an accident, he won't let that deter him from his earlier goal. He takes the pictures out of his car and walks them over to Melissa's house. He carries them over the back fence and into the treehouse, where Henry is sure to see them.

The remaining participants saw a conclusion to the story where Bob had no biological – let alone physical – body at all:

[DISEMBODIED] On the drive over however, Bob is in a fatal car accident and is killed instantly. Bob emerges from his dead body as a ghost. He now has no form at all—no head, no legs, no

arms. Instead, he is something like an invisible force or a spiritual presence. Though he has no limbs with which to touch physical objects, Bob can make objects move without touching them, by floating them through the air. But even though Bob is a ghost, he won't let that deter him from his earlier goal. He causes the pictures to rise out of his car and to float towards Melissa's house. He moves them over the back fence and into the treehouse, where Henry is sure to see them.

All participants were then asked to rate their level of agreement with the following three statements regarding what Bob both felt and believed at the end of the story:

Belief. As Bob moves the pictures into place, he believes Henry will find them in the treehouse after school.

Feel Anger. As Bob moves the pictures into place, he feels angry at Melissa for beginning a new relationship.

Feel Happiness. As Bob moves the pictures into place, he feels happy for Melissa for beginning a new relationship.

Responses were collected on the same seven-item scale anchored with positive and negative agreement terms designed to measure people's willingness to attribute these intentional states (*Belief*) and experiential states (*Feel Anger* and *Feel Happiness*) to Bob.

3.1.2. Results and Discussion

We made three main predictions. A large body of prior empirical work has demonstrated that the behavior displayed by an entity is a crucial factor that cues attributions of intentional states to that en-

tity (Heider & Simmel 1944). So our first prediction was that given Bob's behaviors in the story, participants would signal high levels of agreement with *Belief* across all conditions in the experiment. Second, we predicted that functional information would have a large impact on phenomenal state attribution, whereby people would signal greater agreement with *Feel Anger* in the ANGER condition, and greater agreement with *Feel Happiness* in the HAPPY condition. And lastly, our third prediction was that embodiment would play little to no role in cuing phenomenal state ascription.

All three predictions were borne out. We found that people strongly agreed that Bob believes Henry will find the pictures in the treehouse after school.¹¹ Second, there were large effects for function on the way people attributed emotional states to Bob—differences in Bob's non-phenomenal mental states (i.e., his goals) made a big difference in the phenomenal states that were attributed to Bob. And third, emotional state attributions appeared completely unaffected by whether or not Bob was embodied.¹² These results are

¹¹ *Belief* Disembodied-Anger (M = 6.39, SD = 1.05), Disembodied-Happy (M = 6.42, SD = 0.81), Embodied-Anger (M = 6.54, SD = 0.70), Embodied-Happy (M = 6.17, SD = 1.23).

¹² A 2 (Embodiment) x 2 (Function) MANOVA was used to compare the influence that body salience and function had on the intentional and experiential states: belief, anger and happiness. The multivariate result was significant only for function, Pillai's Trace = 0.48, F = 41.36, df = (3,137), p < 0.001. The univariate F tests showed there was a significant difference between attributions of *Feel Anger* F = 97.79, df = (1,142), p = < 0.001, and *Feel Happiness* F = 101.16, df = (1,142), p = < 0.001 with respect to function. No differences were detected for *Belief*. No main or interaction effects were detected for *Belief*, *Feel Anger*, or *Feel Happiness* with respect to embodiment. We also conducted an additional alternative analysis. Mann-

represented in the figure below:




Figure 1. Mean agreement with mental state attribution in each condition, grouped by mental state probe. All scales ran 1-7. Error bars +/- SE.

Recall that if the embodiment hypothesis is correct in claiming that a crucial factor that cues phenomenal state attribution is whether or not an entity has a certain kind of biological body, then participants should be more likely to disagree that Bob feels anger or happiness when he exists only as an “invisible force or a spiritual presence” as compared to when his biological body is made salient. But what we saw was that whether or not Bob had or lacked a physical body in the various conditions of the experiment (Fig.1: solid vs. patterned

Whitney U test also detected significant differences in *Feel Angry* within both Embodied $U(70) = 178.00$, $Z = -5.29$, $p < 0.001$, and Disembodied $U(71) = 191.00$, $Z = -5.26$, $p < 0.001$ conditions by function.

bars) seemed to play no role in people's willingness to ascribe these mental states to Bob at the end of the story.

On the other hand, functional cues (Fig.1: dark vs. light bars) played a very large role in whether or not people agreed that Bob felt happy or angry. When Bob's goal was to "cause Henry to hate his mom" people were much more likely to agree that Bob feels angry as he places the pictures rather than happy for Melissa for beginning a new relationship. In fact, they attributed this experiential state at roughly the same level as they attributed the intentional state about belief.¹³ Lastly, we observe the opposite pattern when Bob wishes to "make Henry incredibly happy." In such conditions, people were much more likely to agree that Bob feels happy rather than angry with Melissa for beginning a new relationship.¹⁴

3.2. Experiment 2: Disembodied Ghost Lacking Intentional Objects

The findings from Experiment 1 begin to motivate the following conclusions. First, when it comes to entities like disembodied souls, having or lacking a human biological body is not utilized as an important cue when attributing phenomenal consciousness. In fact, this information seemed to play no role in people's judgments. Second, participants strongly agreed by comparison that certain en-

¹³ *Feel Anger* Disembodied-Anger (M = 6.11, SD = 1.19), Disembodied-Happy (M = 3.58, SD = 1.82), Embodied-Anger (M = 6.06, SD = 0.91), Embodied-Happy (M = 3.75, SD = 1.71).

¹⁴ *Feel Happiness* Disembodied-Anger (M = 1.92, SD = 1.65), Disembodied-Happy (M = 4.31, SD = 1.77), Embodied-Anger (M = 1.43, SD = 0.70), Embodied-Happy (M = 4.14, SD = 1.68).

tities—embodied or not—can have experiential states when provided with appropriate functional information.

Regarding the first point however, one immediate objection surfaces. Prior work in experimental philosophy of mind has suggested that participants' agreement with phenomenal state attributions are highly sensitive to the amount of contextual information given within experimental probes. Specifically, Arico (2010) found that attributions of phenomenal states to groups that specified an intentional object (e.g., "Canada's Travel Bureau is experiencing a sudden urge *to pursue internet advertising*") were deemed significantly more acceptable than attributions of phenomenal states to groups that lacked an intentional object (e.g. "Canada's Travel Bureau is experiencing a sudden urge"). In fact, people were much less likely to agree that groups could have a series of phenomenal states when intentional clauses were absent. Thus Arico suggested that that the inclusion of an intentional object in experimental probes provides contextual information that can bias phenomenal ascriptions.

Perhaps a similar effect could explain the high ascriptions of phenomenal states to disembodied entities in Experiment 1. It could be that participants attributed emotional states to Bob because the probes that were used included intentional clauses, (e.g. "he feels angry *at Melissa for beginning a new relationship*" vs. "he feels angry"). These clauses might have served to bias disembodied phenomenal state attributions. We conducted our second experi-

ment to rule out this possibility of bias in the probe design.

3.2.1. Methods

Participants in Experiment 2 (N=147, 53 female, median age = 32) were presented with the same stimulus material combinations as participants in Experiment 1. However after seeing the vignettes, they were asked to rate their agreement with the following three sentences.¹⁵ These sentences were adjusted to account for the worries above by removing the intentional objects from the probe, thereby limiting the potentially biasing contextual information presented:

Intention. Bob intends to move the pictures into place.

Feel Anger No Object. As he moves the pictures into place, Bob feels angry.

Feel Happiness No Object. As he moves the pictures into place, Bob feels happy.

Responses were collected on the same seven-item scale anchored with positive and negative agreement terms.

3.2.2. Results and Discussion

We made two predictions. Our first prediction was that we would replicate each of the results uncovered in Experiment 1. Our second prediction was that the absence of the intentional phrases and potentially biasing contextual information in the phenomenal state probes in Experiment 2 (“he feels angry” vs. “he feels angry at

¹⁵ For the sake of uniformity in removing intentional objects, we switched the intentional state tested in Experiment 2 from believes to intends.

Melissa for beginning a new relationship”) would not result in lower rates of phenomenal state ascription to the disembodied entities in the story.

Both of these predictions were borne out. First, Experiment 2 replicated each effect found in Experiment 1.¹⁶ Once again, participants overwhelmingly attributed *Intention* across the board.¹⁷ Functional cues continued to play a major role in people’s judgments. Participants were much more likely to agree with *Feel Anger No Object* in the ANGER condition, and *Feel Happiness No Object* in the HAPPY condition.¹⁸ And lastly, embodiment again seemed to play no role in cuing phenomenal state ascription. Responses in EMBODIMENT and DISEMBODIMENT were nearly indistinguishable. These results can be seen in Figure 2:

¹⁶ A 2 (Embodiment) x 2 (Function) MANOVA was used to compare the influence that embodiment and function had on the intentional and experiential states: intentionality, anger, and happiness. The multivariate result was significant only for function, Pillai’s Trace = 0.308, $F = 20.94$, $df = (3,141)$, $p < 0.001$. The univariate F tests showed there was a significant difference between attributions of *Feel Anger* $F = 60.54$, $df = (1,146)$, $p = < 0.001$, and *Feel Happiness* $F = 14.34$, $df = (1,146)$, $p = < 0.001$ with respect to function. No differences were detected for *Intention*. And, no main or interaction effects were detected for *Intention*, *Feel Anger No Object*, or *Feel Happiness No Object* with respect to embodiment.

¹⁷ *Intention* Disembodied-Anger (M = 6.50, SD = 0.83), Disembodied-Happy (M = 6.62, SD = 0.49), Embodied-Anger (M = 6.46, SD = 0.56), Embodied-Happy (M = 6.38, SD = 0.59).

¹⁸ *Feel Anger No Object* Disembodied-Anger (M = 5.50, SD = 1.33), Disembodied-Happy (M = 3.49, SD = 1.88), Embodied-Anger (M = 5.46, SD = 1.44), Embodied-Happy (M = 3.41, SD = 1.62). *Feel Happiness No Object* Disembodied-Anger (M = 4.18, SD = 1.72), Disembodied-Happy (M = 5.27, SD = 1.28), Embodied-Anger (M = 3.97, SD = 1.76), Embodied-Happy (M = 4.84, SD = 1.44).

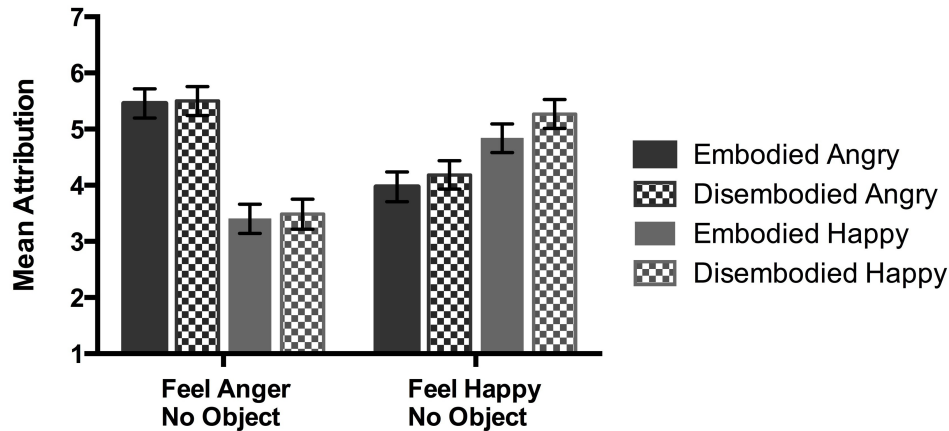


Figure 2. Mean agreement with mental state attribution in each condition grouped by mental state probe. All scales ran 1-7. Error bars +/- SE.

Thus Experiment 2 suggests that Arico's results for contextual information bias for group ascriptions do not extend to phenomenal state ascriptions to disembodied entities such as ghosts. In fact, when the intentional object of the phenomenal state probes are dropped in Experiment 2, we again see a clear demonstration of the role that functional information is playing in people's judgments to these entities.

3.3. Experiment 3: Eternally Disembodied Spirits

Taken together, Experiments 1-2 directly challenge the embodiment hypothesis. But one worry about the entities in these experiments is

that participants might be conceiving of them as *nearly* embodied. After all, Bob did not always lack a unified biological body; he was until very recently a normal human being. So perhaps temporal proximity to unified biological embodiment affects people's judgments about the states they attribute to Bob. It could be that there are specific norms related to the genre of ghost stories such that the ghosts of the recently deceased are attributed phenomenal states because they recently possessed human bodies.¹⁹ To rule out these possibilities we conducted Experiment 3 to see if we could replicate the previous results for entities with no temporal proximity to being normal human beings. For this, we turn to spirits that have never been human, or are *eternally* disembodied.

3.3.1. Methods

Experiment 3 mirrored the same between-subjects multifactorial design in Experiments 1-2. Participants (N=118, 41 female, median age = 30) were presented with cases designed to study the effect of embodiment and functional cues on mental state attribution, this time using an entity that was more purely disembodied. For roughly half of the participants, the story started like this:

[*EMBODIED*] Fintan is a very private person. He has little connection with the outside world—no computer, no phone, no car. Instead, he hunts or grows his own food with his bare

¹⁹ We thank Aaron Meskin for discussion on this point. For more on this point about the kinds of bodies ghosts might be assumed to have see discussion in Section 4 below.

hands. Though he has no money with which to buy tools, Fintan can make many useful objects with the things he finds around him.

The other half saw a story that began as follows:

[DISEMBODIED] Fintan is a nature spirit. He has no form at all —no head, no legs, no arms. Instead, he has always existed as a kind of invisible force or a spiritual presence. Though he has no limbs with which to touch physical objects, Fintan can make objects move without touching them, by floating them through the air.

Both groups then saw the story continue:

For many years, Fintan has lived in Dirk's Wood beside the Mangahala River. He values the beautiful crystal waters and quiet solitude of the Mangahala above everything else. Recently however, construction has started on the Mangahala Golf Course and Retirement Community. Loggers have begun cutting down segments of Dirk's Wood to accommodate the project, polluting the entire area. Fintan decides that the only way to stop the destruction of his home is to cause their trucks and chainsaws to break in any way he can. And when the loggers bring in more equipment, he breaks that too.

Lastly, to manipulate the functional information specified, participants saw one of two conclusions to the story:

[SAD] But the construction company won't give in. Realizing that there is nothing he can do to stop the loggers, Fintan leaves Dirk's Wood. He must now find another place to call home.

[HAPPY] Eventually the construction company gives in. Realizing that he has stopped the loggers, Fintan returns to Dirk's Wood. The place he calls home is now safe.

All participants were then asked the following three questions:

Comprehension. In the story above, Fintan is: [A human being/A nature spirit with no physical body]

Feel Sadness. At the end of the story, Fintan feels sad.

Feel Happiness. At the end of the story, Fintan feels happy.


Phenomenal state ascription was collected on the same seven-item agreement scale used in Experiments 1-2.

3.3.2. Results and Discussion

We made three predictions in Experiment 3. First, we predicted a strong effect for function, whereby people will be much more likely to agree with *Feel Sadness* rather than *Feel Happiness* for SAD, and *Feel Happiness* rather than *Feel Sadness* for HAPPY. Second, we predicted that embodiment would continue to play no role in people's judgments, even when the object of attribution, Fintan, is a nature spirit that has never occupied a physical body. And thirdly, we predicted that we would replicate the earlier finding in Experiment 2, that the absence of intentional object clauses does not preclude phenomenal state attribution.

Again all of these predictions were borne out. We found that

function made a very large difference to phenomenal state attribution in these cases.²⁰ People only attributed *Feel Happiness* or *Feel Sadness* in HAPPY and SAD, respectively.²¹ Attribution between EMBODIED and DISEMBODIED conditions was indistinguishable. And lastly, these results again persisted despite using phenomenal state probes lacking intentional objects. These findings are displayed in Figure 3 below:



²⁰ A 2 (Embodiment) x 2 (Function) MANOVA was used to compare the influence that embodiment and function had on the experiential states anger and happiness. The multivariate result was significant only for function, Pillai's Trace = 0.838, $F = 284.96$, $df = (2,110)$, $p < 0.001$. The univariate F tests showed there was a significant difference between attributions of *Feel Sadness* $F = 374.93$, $df = (1,114)$, $p = < 0.001$, and *Feel Happiness* $F = 542.54$, $df = (1,114)$, $p = < 0.001$ with respect to function. No main or interaction effects were detected for *Feel Sadness* or *Feel Happiness* with respect to embodiment. Three participants were removed for failing Comprehension.

²¹ *Feel Sadness* Disembodied-Sad ($M = 6.00$, $SD = 1.29$), Disembodied-Happy ($M = 2.11$, $SD = 1.40$), Embodied-Sad ($M = 6.14$, $SD = 0.79$), Embodied-Happy ($M = 1.89$, $SD = 0.92$). *Feel Happiness* Disembodied-Sad ($M = 1.57$, $SD = 0.97$), Disembodied-Happy ($M = 5.64$, $SD = 1.19$), Embodied-Sad ($M = 1.66$, $SD = 0.77$), Embodied-Happy ($M = 5.96$, $SD = 0.88$).

Figure 3. Mean agreement with mental state attribution in each condition grouped by mental state probe. All scales ran 1-7. Error bars +/- SE.

3.4. Experiment 4: Eternally Disembodied Spirits—An Alternative Measure

One worry about Experiments 1-3 is that they all use the same basic technique for collecting phenomenal state attributions, in which participants were asked two questions about states of opposite valence. But perhaps presenting these two questions together created some undue pressure to ascribe phenomenal states. With this worry in mind, we conducted Experiment 4 using a different measure for state attribution based on confidence judgments.

3.4.1. Methods

Participants in Experiment 4 (N=120, 37 female, median age = 28) were presented with the same stimulus materials as Experiment 3. However after seeing the materials, they were asked a different set of questions:

Attitude Ascription. Which do you think best describes Fintan at the end of the story? [Fintan feels sad/Fintan feels happy]

Attitude Confidence. How confident are you with the answer you gave to the previous question?

Participants answered *Attitude Ascription* with dichotomous answer choices above. They answered *Attitude Confidence* on a seven-

item scale where “1” was anchored with “Not at all Confident” and “7” was anchored with “Extremely Confident”. *Attitude Ascription* was then recoded (“-1” for feels sad, and “+1” for feels happy) and multiplied by *Attitude Confidence* to create a combined ascription/confidence score (ranging from -7 to +7) for each of the entities in the various combinations of cases.

3.4.2. Results and Discussion

We predicted that this alternative measuring technique in Experiment 4 would still result in the same basic findings as seen in Experiment 3. And that is exactly what we found. Participants still relied on functional information as the crucial cue for ascribing phenomenal states—with total indifference to embodiment.²² Mean combined scores (Attribution x Confidence Rating) were significantly lower when Fintan fails to save his home, and significantly higher when Fintan succeeds in defeating the loggers.²³ These results are shown in Figure 4:

²² A 2 (Function) x 2 (Embodiment) between-subjects analysis of variance reveals a main effect for the factor of Function, $F(1, 119) = 839.08$, $p < 0.001$. No other effects were detected.

²³ Combined Measure Disembodied-Sad ($M = -6.41$, $SD = 1.05$), Disembodied-Happy ($M = 5.34$, $SD = 2.47$), Embodied-Sad ($M = -6.60$, $SD = 0.81$), Embodied-Happy ($M = 4.81$, $SD = 3.25$).

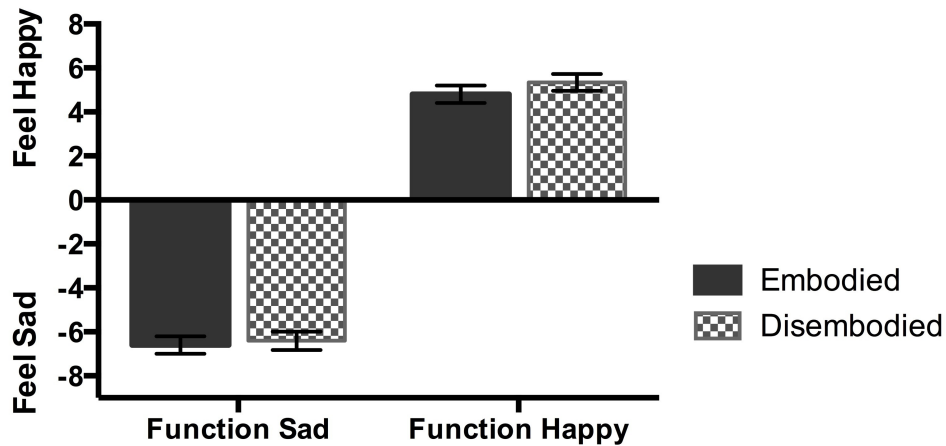


Figure 4. Mean combined score (Attitude Ascription x Attitude Confidence) for each type of entity grouped by function. Scores run from (-7) to (7). Error bars +/- SE.

3.5. Experiment 5: Spirits, Groups, and Humans—Explicit Emotional Comparisons

The previous experiments appear to demonstrate that people often ascribe emotional states to disembodied entities without hesitation—so long as the appropriate functional cues are present. But then again, how can we be sure that participants are applying the phrase “feels sad” to a disembodied entity as they would to a normal human? Perhaps people merely *say* that the spirit feels sad, but *mean* something different than what they mean when they say a human being is sad. In other words, they might not literally attribute the

state of sadness to a spirit in the same way they do to a normal human being. In that case, let's say they make an *anti-realist* ascription.

To ensure that participants are literally ascribing phenomenal states in both cases, we need evidence that when people ascribe emotional states to ghosts and spirits, they mean to attribute the same emotional states they attribute to other human beings when they make similar ascriptions. In other words, we need evidence of *realist* ascriptions. In our fifth study, we set out to provide such evidence.

As recent experimental work on quantity implicatures (in addition to other work in experimental pragmatics) demonstrates, it is often very difficult to experimentally uncover what people mean by (or how they interpret) particular sentences.²⁴ However, our task is at least somewhat less daunting since we do not need to uncover what people ultimately mean when they say a spirit is sad. We simply need to demonstrate that people generally mean the same thing by “sad” when they say, for instance, “a spirit is sad” as they do when they say “a person is sad”. There may be numerous ways of examining this question. But one straightforward way is just to ask people to evaluate their mental state ascriptions comparatively. In other words, we could simply ask those who ascribed emotional states to the spirits how similar the emotional states they meant to attribute were to the emotional states they would attribute to a nor-

²⁴ See Noveck and Reboul (2008) for a useful review.

mal person.

Of course, we would expect some variance in individual responses to this question, so we would need to compare responses to a similar question asked of those who more or less agreed with emotional state ascriptions to the human character in our stories as well. And since we were predicting no difference between people's interpretations of emotion words for the spirit or the man, we would also need some other entity to serve as a control, some entity to which people are willing to ascribe emotional states at the verbal level, but to which they do not really mean to attribute exactly the same emotional states they attribute to normal persons.

For this, we turn again to prior work in the experimental philosophy of mind on group ascriptions. Specifically, Phelan et al. (2012) found that people often offer anti-realist phenomenal state ascriptions to group entities (e.g. the Boeing Corporation).²⁵ Thus group entities seem like the perfect control to use in Experiment 5 when checking for realist ascriptions. Recall the Microsoft Corporation example in Section 1. According to the embodiment hypothesis, people should be thinking about group agents in the same way that they are thinking about disembodied spirits. That is, people should be hesitant to make realist phenomenal ascriptions to both sorts of entities because they lack the right kind of body. So in what follows,

²⁵ Phelan et al. (2012) argue that phenomenal state ascriptions to groups are often *distributivist*, or that people ascribe states to individual group members rather than to the group as a whole over and above its members.

we reexamine participants' judgments by asking for explicit comparisons between different emotional states of humans, disembodied spirits, and group entities.

3.5.1. Methods

Participants (N=194, 75 female, median age = 26) read vignettes similar to those used in Experiments 3-4. Each vignette began with the introduction of a protagonist that was either a spirit, a human, or a group:

[SPIRIT] Fintan is a nature spirit who strives to protect local forests and rivers. He has no form at all—no head, no legs, no arms. Instead, he has always existed as a kind of invisible force or a spiritual presence. Though he has no limbs with which to touch physical objects, Fintan can make objects move without touching them, by floating them through the air. He uses his spiritual abilities to bring an active approach to nature preservation.

[HUMAN] Fintan is an individual who strives to protect local forests and rivers. Through hard work and tireless efforts, Fintan works to protect natural areas from development. Though he has little money with which to support his cause, Fintan exploits his own significant technical skills to bring an active approach to nature preservation.

[GROUP] FINTAN is an organization set up to protect local forests and rivers. Through charitable donations and the efforts of group members, FINTAN works to protect natural areas from

development. Not only does FINTAN support conservation legislation, it also exploits the technical skills of members to bring an active approach to nature preservation.

All participants then read a short description of the character's struggle against a development project, for the spirit it read as follows (with only necessary changes to the character made across other vignettes):

For many years, Fintan has worked to protect Dirk's Wood beside the Mangahala River. The spirit values the beautiful crystal waters and quite solitude of the Mangahala above everything else. Recently however, construction has started on the Mangahala Golf Course and Retirement Community. Loggers have begun cutting down segments of Dirk's Wood to accommodate the project, polluting the entire area. After an extended struggle, Fintan decides that the only way to stop the destruction of the woods is to cause the loggers' trucks and chainsaws to break in any way possible. When the loggers bring in more equipment, Fintan breaks those too.

Lastly, participants were presented with one of two possible endings to the story, where Fintan is either successful or unsuccessful at thwarting the logger's effort:

[HAPPY] Eventually the construction company gives in. Realizing that the loggers have been stopped, Fintan celebrates the preservation of Dirk's Wood.

[SAD] But the construction company won't give in. Realizing that nothing can be done to stop the loggers, Fintan gives up on Dirk's Wood.

All participants were then asked the following two questions:

Comprehension. In the story above, Fintan is: [A group/A nature spirit with no physical body/A human being]

Emotional State Attribution. At the end of the story, Fintan feels [sad/happy].

Finally, those participants who answered *Emotional State Attribution* with "somewhat agree", "agree", or "strongly agree" proceeded to the additional follow-up question designed to measure realist ascriptions:

Comparison. Consider [Fintan the spirit/Fintan the human being/FINTAN the group (over and above the people that constitute it)]. When you say that this [spirit/man/group] feels [sad/happy], how similar is the feeling of [sadness/happiness] to that of a normal person?

Participants responded to this question by rating their level of agreement on a seven-item scale, running from 'Not at all Similar' to 'Exactly the Same.'

3.5.2. Results and Discussion

Given the results of Experiment 1-4, we expected high *Emotional state attribution* in HAPPY and SAD. And that is exactly what we found. The vast majority of participants (over 80% per condition) signalled at least some agreement with *Emotional State Attribution*.

tion.²⁶

We now move on to an analysis of *Comparison* for those participants (N=174) agreeing with *Emotional State Attribution*. If the ascriptions people make to the spirit are realist ascriptions—that is, if people mean the same thing when they attribute emotional states to spirits as they do when they attribute emotional states to people—then we would expect (1) high *Comparison* scores in both HUMAN and SPIRIT (2) with no significant differences between scores in these conditions. And since prior work has suggested that phenomenal state ascriptions to groups are non-realist ascriptions, we would expect (3) significantly lower *Comparison* scores in GROUP than in both HUMAN and SPIRIT.

Again, this is exactly what we found. Despite high scores, there was no significant difference in *Comparison* between HUMAN and SPIRIT. Participants indicated that they generally mean the same thing by “feeling happy” or “feeling sad” when directed towards a disembodied spirit or a human being. And consistent with prior findings, we also found that *Comparison* judgments were significantly lower in GROUP than they were for both HUMAN and

²⁶ Percent agreement with the emotional state ascription per condition: ‘Sad Spirit’ = 87.9%; ‘Sad Group’ = 81.8%; ‘Sad Man’ = 93.9%; ‘Happy Spirit’ = 93.9%; ‘Happy Group’ = 90.6%; ‘Happy Man’ = 96.7%. Overall, people were somewhat more likely to ascribe happiness (M=6.22) than sadness (M=5.79). A 2 (Emotional state) x 3 (Entity Type) between-subjects analysis of variance reveals a main effect for the factor emotional state, $F(1, 193) = 839.08, p < 0.05$. No other effects were detected. We set this result aside.

SPIRIT.²⁷ Results for Comparison are shown in Figure 5 below:



Figure 5. Mean Comparison judgment for emotional attribution grouped by entity type. All scales ran from 1 to 7. Error bars +/- SE.

These results suggest that people think of the emotional states they attribute to disembodied entities in the same way as they think of the emotional states they attribute to human beings. In other words, this is evidence that they think these states are similar to the

²⁷ A 2 (Emotional State) x 3 (Entity Type) between-subjects analysis of variance reveals a main effect for the factor Entity Type, $F(2, 173) = 10.29$ $p < 0.001$. No other effects were detected. A Tukey HSD test revealed significant differences for people's interpretations of emotional state attributions between GROUP and both SPIRIT ($p = 0.001$) and HUMAN ($p = 0.001$). However, no significant difference emerged for HUMAN and SPIRIT ($p = 0.897$). Six participants were removed for failing *Comprehension*.

emotional states of normal people. And they think these states are somewhat dissimilar from the emotional states they attribute to groups. Lastly, recall that the embodiment hypothesis predicted that people would be thinking about groups and disembodied spirits in the same way. But it turns out that we see very different results when using our realist measure for comparing ascriptions to these two types of entities.

4. Feeling Beyond Embodiment

Our experiments suggest that people are perfectly willing to ascribe emotional states to disembodied entities (ghosts and spirits). Though we think more experiments need to be conducted pursuing the question of realist ascription, we think that these results are a promising first step towards the conclusion that findings across Experiments 1-5 constitute strong evidence against the embodiment view. It appears that people really do think that under the right conditions, disembodied entities can have the same kinds of emotional states as human beings. What's more, the data from study 5 suggest that people think of emotional state ascriptions to disembodied entities *in the same way* as they think of emotional state ascriptions to human beings.

Of course, even though participants explicitly state that entities like Fintan are disembodied, it could be that there are specific cultural or social norms which nonetheless suggest that all spirits

occupy a location, and thus must possess a body in some indeterminate or minimal sense.²⁸ Indeed there probably is such a sense in which spirits have bodies, much like there is some mitigated sense in which group entities like Microsoft have bodies. We would only point out that the crucial question – and perhaps the feature that attracted many to the embodiment hypothesis in the first place – was whether or not phenomenal ascriptions are cued in light of possessing a *unified biological body* like our own. It is unclear whether the minimal or indeterminate sense in which ghosts might be assumed to have bodies meets with these criteria.

We should also point out that while we found strong evidence for phenomenal state ascriptions to entities lacking unified biological bodies, embodiment could still have a *relative* impact on ascription. In other words, it's possible that people attribute more, or will be more likely to attribute certain phenomenal states or mental capacities to entities as considerations about the body become *more* salient.²⁹ While this continues to be a possibility note that in our experiments we found extremely similar rates of ascription between embodied and disembodied entities. If embodiment made any kind of minimal incremental difference on phenomenal ascription, we did not detect it across our experiments. In contrast, we present strong evidence for the distinct and central role that function plays in ordinary judgments.

²⁸ We thank Joshua Weisberg for discussion on this point.

²⁹ We thank Shaun Nichols for discussion on this point.

These results corroborate previous research by Buckwalter & Phelan on the important—*otherworldly* even—role that functional information has on the ascription of phenomenal states to diverse sorts of entities. Completely independently of any influence for body, information about perceptual stimuli, behavioral responses, and other mental states is accompanied by strong attribution of phenomenal states, including, as discussed here, emotional states. One straightforward explanation of this fact is that folk psychology actually identifies phenomenal states with functional roles. This is an interesting question to be pursued in future research on phenomenality and functional role. However for our current purposes, what we find striking is that people's judgments were highly sensitive to functional role in exactly the same manner for both entities with or without unified biological bodies. We conclude that when it comes to the psychological factors that cue people's actual attributions of phenomenal states to ghosts, perhaps the only *apparition* here is the embodiment hypothesis itself.

Returning now to the “Nation of China” thought experiment, it could be that Block's basic insight about what is ultimately required for phenomenal consciousness is still more or less on the right track. After all, the experiments we conducted only speak to the principles of folk psychology that guide ordinary ascriptions of phenomenal states. They don't rule out the metaphysical possibility that cognition requires some sort of embodiment. While this remains a possibility, we would only note that part of the argument

for this metaphysical picture of the mind was motivated by the intuition that “China-Brain” does not have mental states in the first place. But if our results for disembodied entities are shown to be sufficiently general, this intuition may not be widely shared.

REFERENCES

- Arico, A. 2010. “Folk Psychology, Consciousness, and Context Effects.” *Review of Philosophy and Psychology*, 1(3): 371-393.
- Bering, J. M. 2002. “Intuitive conceptions of dead agents’ minds: The natural foundations of afterlife beliefs as phenomenological boundary,” *Journal of Cognition and Culture*, 2, 263-308.
- Bering, J. M. 2011. *The God Instinct: The Psychology of Souls, Destiny, and the Meaning of Life*. London: Nicholas Brealey Publishing.
- Bering, J. and Bjorklund, D. 2004. “The Natural Emergence of Afterlife Reasoning as a Developmental Regularity.” *Developmental Psychology*, 40, 217-33.
- Bering, J., Hernández-Blasi, C., and Bjorklund, D. 2005. “The Development of “Afterlife” Beliefs in Religiously and Secularly Schooled Children,” *British Journal of Developmental Psychology*.

- Block, N. 1978. "Troubles with Functionalism" in W. Savage, ed., *Perception and Cognition: Issues in the Foundations of Psychology*, Minnesota Studies in the Philosophy of Science, vol. IX, pp. 261-326.
- Buckwalter, W., and Phelan, M. 2012. Function and Feeling Machines: A Defense of the Philosophical Conception of Subjective Experience. *Philosophical Studies*. DOI 10.1007/s11098-012-0039-9.
- Dennett, D. C. 1991. *Consciousness Explained*. Boston: Little, Brown and Company
- Gray, K., & Wegner, D. M. 2010. Blaming God for Our Pain: Human Suffering and the Divine Mind. *Pers Soc Psychol Rev*, 7-16.
- Gray, H., Gray, K., and Wegner, D. M. 2007. "Dimensions of mind perception." *Science*, 315, 619.
- Gray, H., Knobe, J., Sheskin, M., Bloom, P., & Barrett, L., B. 2011. "More Than a Body: Mind Perception and the Nature of Objectification" *Journal of Personality and Social Psychology*, 101, 6, 1207–1220.
- Heider, F. and Simmel M. 1944. "An experimental study of apparent behavior." *American Journal of Psychology* 57, 243–259.
- Knobe, J. 2008. "Can a robot, an insect, or God be aware?" *Scientific American: Mind*.

- Knobe, J. and Prinz, J. 2008. "Intuitions about consciousness: Experimental studies." *Phenomenology and the Cognitive Sciences*, 7, 67–83.
- Nichols, S. 2007. "Imagination and Immortality: thinking of me." *Synthese*, 159(2): 216-233.
- Noveck, I. A. and Reboul, A. 2008. "Experimental Pragmatics: A Gricean turn in the study of language." *Trends in Cognitive Science*, 12(11), 425-31.
- Phelan, M., Arico, A., and Nichols, S. 2012. "Thinking things and feeling things: On an alleged discontinuity in the folk metaphysics of mind." *Phenomenology and the Cognitive Sciences*. DOI: 10.1007/s11097-012-9278-7
- Phelan, M. and Buckwalter, W. Forthcoming. "Analytic Functionalism and Mental State Attribution." *Philosophical Topics*.
- Searle, J. 1980. "Minds, Brains, and Programs," *Behavioral and Brain Sciences*, (3): 417-458.
- Sytsma, J., and Machery, E. (2009). "How to study folk intuitions about phenomenal consciousness," *Philosophical Psychology* 22 (1):21 – 35.
- Sytsma, J., and Machery, E. (2010). "Two conceptions of subjective experience," *Philosophical Studies* 151 (2):299-327.