

Attributions of Consciousness¹

Justin Sytsma

Abstract: Many philosophers and brain scientists hold that explaining consciousness is one of the major outstanding problems facing modern science today. One type of consciousness in particular—phenomenal consciousness—is thought to be especially problematic. The reasons given for believing that this phenomenon exists in the first place, however, often hinge on the claim that its existence is simply obvious in ordinary perceptual experience. Such claims motivate the study of people’s intuitions about consciousness. In recent years a number of researchers in experimental philosophy of mind have begun to shed light on this area, investigating how people understand and attribute those mental states that have been thought to be phenomenally conscious. In this article I will discuss the philosophical concept of phenomenal consciousness and detail the work that has been done on the question of whether lay people have this concept.

It is approaching a cliché to claim that explaining consciousness is one of the major outstanding scientific problems of the 21st century. Over the past 20 years, a host of prominent brain scientists have lined up to note the puzzle and urge that researchers take it seriously. To give but one example, co-discover of DNA Francis Crick and neuroscientist Christof Koch (2004, 273) write: “The explanation of consciousness is one of the major unsolved problems of modern science. After several thousand years of speculation, it would be very gratifying to find an answer to it.” Gratifying indeed! After all, it is commonly claimed that consciousness is the most important aspect of our mental lives.

Not surprisingly, many scientists have taken up the challenge of explaining consciousness. As psychologist Steven Pinker writes in *Time* magazine (January 19, 2007), “questions once

¹ To appear in *WIREs Cognitive Science*. I would like to thank Wesley Buckwalter for his insights during the early stages of this project, as well as Simon Keller, Ramon Das, Nicky Drake, Anton Killin, and two anonymous referees for *WIREs Cognitive Science* for their helpful comments on a previous draft of this article.

confined to theological speculations and late-night dorm-room bull sessions are now at the forefront of cognitive neuroscience.” The most puzzling of these questions is the one that philosopher David Chalmers (1995, 1996) has labeled the hard problem of consciousness. Pinker describes the problem as follows:

The Hard Problem... is why it feels like something to have a conscious process going on in one’s head—why there is first-person, subjective experience. Not only does a green thing look different from a red thing, remind us of other green things and inspire us to say, “That’s green” (the Easy Problem), but it also actually looks green: it produces an experience of sheer greenness that isn’t reducible to anything else. As Louis Armstrong said in response to a request to define jazz, “When you got to ask what it is, you never get to know.”²

Put another way, the hard problem of consciousness is to explain what philosophers typically call *phenomenal consciousness*. Of course, phenomenal consciousness is certainly not the only thing that might be meant by the term “consciousness.” The term is used in a number of different senses—it is variously used to indicate that an entity is awake (as opposed to asleep or comatose), sentient, aware of something, or able to report on something, amongst other usages. In this article, however, I will focus on attributions of phenomenal consciousness, leaving to the side interesting questions about how people attribute other types of consciousness (but see Wilkes, 1988; De Brigard, 2010).

While researchers are not in perfect agreement about how the concept of phenomenal consciousness should be understood, the standard line is that a creature is phenomenally conscious just in case it has *phenomenally conscious mental states*, and that a mental state is phenomenally conscious just in case it has *phenomenal qualities* (*qualia* for short). Qualia are typically understood as being those qualities of phenomenally conscious mental states that give them a distinctive feel: It is thought that there is “something it is like” (Nagel, 1974) to be in a

² Pinker borrows this analogy from Ned Block’s classic paper “Troubles with Functionalism.” Block (1978, 281) writes: “You ask: What is it that philosophers have called qualitative states? I answer, only half in jest: As Louis Armstrong said when asked what jazz is, ‘If you got to ask you ain’t never gonna get to know.’”

phenomenally conscious mental state, and this “something” corresponds with the state’s qualia. This is often illustrated by listing off ordinary perceptual episodes in which one is acquainted with certain distinctive qualities. For example, Koch (2004, 1-2) writes:

How do the salty taste and crunchy texture of potato chips, the unmistakable smell of dogs after they have been in the rain, or the feeling of hanging on tiny fingerholds on a cliff a couple of meters above the last secure foothold, emerge from networks of neurons? These sensory qualities, the building blocks of conscious experience, have traditionally been called qualia. The puzzle is, how can a physical system have qualia?

The distinctive qualities that Koch notes (the taste, the texture, the smell, the feeling) are taken to be qualia—they are treated as being qualities *of* the corresponding mental states. In other words, in asserting that these qualities are qualia, they are taken to be mind dependent in the sense that if there were no mental states then there could be no qualia.

A number of researchers, both philosophers and scientists, have thought that phenomenal consciousness exists and that its existence is puzzling, forming a significant standout against the steady march of empirical explanation. Of course, many phenomena have resisted scientific explanation for a period of time, and many continue to do so today. For example, the details of many complex processes—such as the acquisition of syntax by children, for instance—remain largely unknown to us. The supposed puzzle of phenomenal consciousness goes beyond this, however: Many hold that we not only do not know how to fill out the details of an adequate account of phenomenal consciousness, but that we do not even know what such an account should look like in outline. It is this broad recalcitrance to scientific explanation that is intended when phenomenal consciousness is said to pose a distinctively hard problem.

Going the other direction, however, the claim that phenomenal consciousness poses such a difficult puzzle might itself be thought to be puzzling. Flipping the issue on its head, one might take the fact that phenomenal consciousness seems to be so singularly resistant to empirical

explanation to give us reason to doubt that we have adequately delineated the supposed phenomenon. In fact, one might take the claim that phenomenal consciousness poses a hard problem to provide reason to doubt the reality of this supposed phenomenon (see Sytsma, 2009, 2010a, for an exploration of this idea). With this skeptical motivation in place, we might then ask: Why have so many researchers thought that this mysterious phenomenon exists in the first place?

The basic answer is that the existence of phenomenal consciousness is taken to follow from its supposed importance to everyday life. In other words, the existence of phenomenal consciousness is thought to be evident when you pay attention to your own mental life. It is thought to be simply manifest in introspection when one eats a potato chip, or smells a wet dog, or hangs from a cliff a couple of meters from a secure foothold. To borrow a phrase from philosopher Daniel Stoljar (2006, v), phenomenal consciousness is thought to be *phenomenologically obvious*. And this is basically the justification that Chalmers offers in introducing the hard problem of consciousness. The idea is that “there is nothing that we know more intimately than conscious experience” (Chalmers, 1995, 200), as it is “the most central and manifest aspect of our mental lives” (207).

The reason offered for believing that phenomenal conscious exists is similar to what we find behind a number of philosophical puzzles—the central claim seems to rest on how things seem to the researchers (or on the researchers’ intuitions, as it is often put).³ Of course, we might worry that how things seem to these researchers might not reflect what is pretheoretically obvious about their mental lives, but the concepts they have learned and the views that they have arrived at over the course of their extensive training with regard to these matters. Continuing with the skeptical line introduced above, we might then worry that the existence of phenomenal

³ For a discussion of the use of intuitions in philosophy—and some of the worries that have been raised about this practice by experimental philosophers—see Alexander (2012), especially Chapters 1 and 2; see also Sytsma and Livengood (forthcoming), especially Chapter 2.

consciousness is not actually manifest in ordinary experience as many researchers have claimed. And if that is the case, then we might be concerned that far from being “one of the major unsolved problems of modern science,” the existence of phenomenal consciousness is in fact just a pseudo-problem.

Such worries are one motivation for not just studying *consciousness*, but studying people’s *intuitions about consciousness*: Rather than just accept claims that the existence of phenomenal consciousness is manifest in ordinary experience, we can investigate the issue further by exploring how lay people (the “folk” or people without training in philosophy or consciousness studies) tend to think about those mental states that philosophers have typically taken to be phenomenally conscious. For example, if it is truly obvious from first-person experience with such states that they are phenomenally conscious, then we should expect lay people to recognize this fact. But do they?

Perhaps surprisingly given all of the effort that has been put into trying to explain phenomenal consciousness, relatively little work has been put into investigating people’s intuitions about consciousness.⁴ This has begun to change in recent years, however, with a spate of new research in the burgeoning area of experimental philosophy of mind investigating this issue.⁵ I will review this literature in this article, targeting the question of whether lay people employ the concept of phenomenal consciousness. I begin in Section 1, by looking at the evidence for an affirmative answer to this question, focusing on Joshua Knobe and Jesse Prinz’s (2008) article, “Intuitions About Consciousness: Experimental Studies.” I then consider some

⁴ Interestingly, a great deal of work has been done over the past several decades concerning how people understand and attribute other types of mental states. As described in Arico (2010) and Sytsma (2010b), work on theory of mind has largely focused on mental states such as beliefs and desires that are involved in predicting agentive behavior, but are not typically thought to be phenomenally conscious (see, for example, Wellman, 1990; Perner, 1991; Gopnik and Metzliff, 1997; Baron-Cohen, 2001; Malle, 2005; Onishi and Baillargeon, 2005).

⁵ For an accessible look at the broader area of experimental philosophy of mind, see Machery and Sytsma (2011) and Knobe (2008); for a collection of cutting-edge research in this area, see the articles in Sytsma (2014a).

objections that have been raised against this work in Section 2. In Section 3, I flip the table and look at the evidence for a negative answer to the target question, focusing on Justin Sytsma and Edouard Machery's (2010) article, "Two Conceptions of Subjective Experience." I then consider some objections that have been raised against this work in Section 4.

1. The Positive Case for a Lay Concept of Phenomenal Consciousness

While philosophers working on consciousness often call on their own intuitions, few have actually investigated people's intuitions about consciousness. Knobe and Prinz (2008, 67) lament this fact:

This emphasis strikes us as unfortunate. Intuitions about consciousness are fascinating phenomena, amply worthy of study in their own right. The fact that people have the intuitions they do can teach us something valuable about the way people ascribe mental states, the way they think about nonhuman animals, perhaps even the way they make moral judgments.

They then attempt to rectify this situation, carrying out a series of empirical studies to investigate lay people's intuitions about consciousness. Amongst other things, Knobe and Prinz wanted to know whether lay people employ the concept of phenomenal consciousness. Since "phenomenal consciousness" is not everyday terminology, however, and since "consciousness" can be used in many different senses, they couldn't just ask people this question directly. Instead, what Knobe and Prinz did was to look at how lay people ascribe different types of mental states to different entities. What they wanted to know is whether lay people would tend to treat mental states that philosophers classify as phenomenal differently than those that philosophers classify as non-phenomenal. Knobe and Prinz argued that if people did so, then this would be evidence that they employ the concept of phenomenal consciousness: It would indicate that they treat phenomenally

conscious mental states as forming a type, suggesting that they take these states to have something in common, namely that they are all phenomenally conscious.

Since Knobe and Prinz wanted to determine whether lay people would tend to treat phenomenal mental states differently from non-phenomenal mental states, they couldn't ask people about just any type of entity. After all, people might think that the entity was capable of *both* types of states, as we would expect for attributions to a normal adult human. What Knobe and Prinz needed, instead, was an entity that they expected lay people would think of as being incapable of having one of these two types of mental states. They predicted that this would be the case for group agents, such as the corporation Microsoft. Specifically, Knobe and Prinz predicted that lay people would be willing to ascribe non-phenomenal mental states to such an entity, but that they would be unwilling to ascribe phenomenal mental states to it.

This prediction followed from two hypotheses. First, Knobe and Prinz hypothesized that lay people call on the concept of phenomenal consciousness in deciding whether or not to attribute phenomenally conscious mental states to a given entity, recognizing that the entity must be phenomenally conscious to be in such states. Second, they hypothesized that the *physical constitution* of an entity—i.e., the physical make-up of the entity, or what it is made out of and how it is put together—matters for lay attributions of phenomenally conscious mental states in a way that it does not for attributions of other mental states (such as beliefs or desires). Specifically, Knobe and Prinz suspected that *embodiment* matters for judgments about phenomenal consciousness, with lay people being more likely to attribute phenomenally conscious mental states to entities that have an individual body than to entities that are distributed (such as a group agent).⁶

⁶ Although I will focus on Knobe and Prinz's first hypothesis in this article, it is worth noting that doubts have also been raised about their embodiment hypothesis and similar claims found in Knobe (2008) and Gray et al. (2011a). Buckwalter

Combining these two hypotheses, Knobe and Prinz predicted that lay people would be less willing to attribute mental states that philosophers typically consider to be phenomenally conscious to a group agent than they would be to attribute mental states that are not typically considered to be phenomenally conscious. To test this prediction, in their second study they gave participants the ten sentences below, counterbalanced for order:

Acme Corp. has just decided to adopt a new marketing plan.
Acme Corp. wants to change its corporate image.
Acme Corp. intends to release a new product this January.
Acme Corp. believes that its profit margin will soon increase.
Acme Corp. knows that it can never compete with GenCorp in the pharmaceuticals market.
Acme Corp. is experiencing a sudden urge to pursue internet advertising.
Acme Corp. is now experiencing great joy.
Acme Corp. is now vividly imagining a purple square.
Acme Corp. is getting depressed.
Acme Corp. is feeling excruciating pain.

Each of these sentences attributes a different mental states to the group agent Acme Corporation. Participants were then asked to rate each sentence using a 7-point scale that was anchored at 1 with “sounds weird” and at 7 with “sounds natural.”

Knobe and Prinz judged that the first five sentences in the list above involve non-phenomenal mental states, while the last five involve phenomenal mental. As such, they predicted that participants would rate the latter set of sentences as being more weird sounding than the former. And this is exactly what they found, as shown in Figure 1. Most importantly, the mean rating for each of the five sentences attributing a mental state that Knobe and Prinz judged to be non-phenomenal was higher than for each of the five sentences attributing a mental state

and Phelan (2014) note that previous research by Bering and colleagues (Bering, 2002; Bering and Bjorklund, 2004; Bering et al., 2005) on attributions of psychological functioning after biological death suggest against the embodiment hypothesis, and present the results of five new studies that directly challenge the hypothesis. Across their studies, Buckwalter and Phelan found that lay people tended to be willing to attribute emotional states (anger, happiness, sadness) to disembodied ghosts and spirits: “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... suggest that people think of emotional state ascriptions to disembodied entities *in the same way* they think of emotional state ascriptions to human beings.” (67).

that they judged to be phenomenal.⁷ Following the reasoning presented above, Knobe and Prinz took at these findings to suggest that lay people do in fact employ the concept of phenomenal consciousness: They took these results to indicate that lay people treat phenomenally conscious mental states as forming a type, distinguishing them from other mental states, which in turn suggests that they see these states as having something important in common—i.e., that they are phenomenally conscious.




Figure 1: Results of Study 2 in Knobe and Prinz (2008), showing the mean response for each sentence in descending order (1= “sounds weird” and 7=“sounds natural”).

⁷ This study has been replicated by Sytsma and Machery (2009) using the same sentences as Knobe and Prinz. In addition, similar results were found by Huebner, Bruno, and Sarkissian (2010) using a set of 20 sentences attributing mental states to a range of group agents (Denmark, Destiny’s Child, the Catholic Church, Sony, etc.). In their cross-cultural study, they found that participants from the University of North Carolina rated each of the sentences that Huebner and colleagues classified as non-phenomenal as sounding more natural than each of the sentences that they classified as phenomenal. A corresponding trend was found for participants from the Chinese University of Hong Kong, with mean ratings for the sentences classified as non-phenomenal being higher than for the sentences classified as phenomenal. However, these participants rated some phenomenal states as sounding more natural than some non-phenomenal states.

1.1 Related Findings

Before turning to criticisms of Knobe and Prinz's work, it is worth noting that a good deal of research has been done on the closely related thesis that there are two distinct dimensions of mind perception (Gray et al., 2007, 2011b, 2012a, 2012b; Gray and Wegner, 2009, 2010, 2011a, 2011b; Gray and Schein, 2012; Waytz et al., 2010; see Theriault and Young, 2014, for an overview). In their seminal article arguing for this thesis, Gray, Gray, and Wegner (2007) presented participants with 78 pair-wise comparisons of 13 characters (such as a 5-year-old girl, an adult man, a wild chimpanzee, God) and solicited ratings for one of 18 mental capacities (such as the capacity to feel pain) or six personal judgments (such as which character they liked more). Using a principal components factor analysis, they found that the mental capacities tested divided into two factors, which they termed Experience and Agency.⁸

Although the name chosen for Gray and colleagues' Experience dimension is suggestive of the concept of phenomenal consciousness, it is not clear that it should be interpreted in this way when we consider the items that comprise it—hunger, fear, pain, pleasure, rage, desire, personality, consciousness, pride, embarrassment, and joy. While some of these items are prototypical examples of phenomenally conscious mental states (such as hunger, pain), others are less clear (such as personality, consciousness, pride). As such, while Gray and colleagues' results are perhaps suggestive of Knobe and Prinz's claim that lay people distinguish between mental

⁸ Gray, Gray, and Wegner (2007) link their two dimensions of mind perception to moral cognition, and specifically the distinction between moral agents (capable of morally right or wrong actions) and moral patients (morally right or wrong actions can be done to them)—Experience being linked to moral patiency, Agency to moral agency. Much of the subsequent literature exploring the two dimensions of mind perception has explored this connection, although Sytsma and Machery (2012a) and Piazza, Landy, and Goodwin (forthcoming) argue that the connections are not quite so neat as this suggests. See Theriault and Young (2014) for an accessible review. Similar connections have been drawn between moral cognition and the distinction between phenomenal and non-phenomenal mental states (see Phelan and Waytz, 2012, for discussion). Thus, Knobe and Prinz (2008, fn7) note that “phenomenal consciousness is specifically relevant to judgments of *moral patiency*... rather than to judgments of *moral agency*... Judgments of moral agency appear to depend more on ascriptions of non-phenomenal states, such as beliefs, desires and intentions.” A similar connection is drawn by Robbins and Jack (2006).

states that are and that are not phenomenally conscious, we should be careful about treating this study as providing direct evidence for that hypothesis (for discussion, see Huebner, 2010, 135–136; Sytsma and Machery, 2010, 302–303; Phelan et al., 2013, 705). This difficulty is nicely illustrated by the work of Haslam et al. (2008), who note that Gray and colleagues’ findings accord with the two-factor model of dehumanization put forward in Haslam (2006), even though these factors do not correspond with the distinction between phenomenal and non-phenomenal mental states.⁹ A similar point can be made concerning the work of Philip Robbins and Anthony Jack (2006; Jack and Robbins, 2012; Jack et al., 2014; Jack, forthcoming), who explicitly equate the Experience dimension with phenomenal consciousness, arguing that in addition to the physical and intentional stances articulated by Daniel Dennett (1981, 1987), people employ a phenomenal stance.

Bryce Huebner (2010) also tackles the question of whether lay people have the concept of phenomenal consciousness, presenting the results of two studies involving ascriptions of mental states to either a human, a robot, or two types of cyborgs (human brain and robot body, robot body and human brain). In each study, participants were given a description of the agent including a picture, then asked whether they agreed or disagreed with two statements, one involving a phenomenal state (“he feels pain if he is injured or damaged in some way,” “he feels happy when he gets what he wants”) and one involving a non-phenomenal state (“he believes that $2+2=4$,” “he believes that triangles have three sides”). As shown in Figure 2, Huebner found that while participants tended to treat the human and the robot similarly with regard to the non-phenomenal states (they were generally willing to attribute belief to both the human and the robot), they tended to treat the human and the robot dissimilarly with regard to the phenomenal

⁹ Likewise for Fiske et al.’s (2002) model of social evaluation, which distinguishes between competence and warmth, as noted by Waytz et al. (2010).

states (they were generally willing to attribute both pain and happiness to the human, but not the robot).




Figure 2: Results from Huebner (2010) showing percentage of participants attributing phenomenal (pain, happiness) versus non-phenomenal (belief) mental states to a human and a robot.

Huebner notes that at first glance, his results might appear to support the view advanced by Knobe and Prinz (2008), in that they “seem to confirm that commonsense psychology does draw a distinction between phenomenal and non-phenomenal states—and this distinction seems to be dependent on the structural properties of an entity in a way that ascriptions of non-phenomenal states are not” (138–139). He goes on to argue that to move from his results to such a conclusion is to move far too quickly, however. Huebner’s basic worry is that while his results are compatible with this conclusion, they are also compatible with the opposite claim—that lay people do *not* employ the concept of phenomenal consciousness. The reason is that pain/happiness can be

distinguished from beliefs in a number of ways beyond their supposed phenomenality. For example, we might distinguish between these states in terms of whether they are typically associated with a hedonic value (Sytsma and Machery, 2010) or simply whether they are thought to be beyond the reach of robots (Fiala, Arico, and Nichols, 2014). What this means is that people might treat robots as being capable of belief and incapable of pain or happiness without possessing or employing the philosophical concept of phenomenal consciousness. As we will see below, a similar worry has been raised about Knobe and Prinz's findings, even though they asked participants about a broader range of mental states than did Huebner.¹⁰

2. Alternative Explanations

At present, the strongest case for the claim that lay people tend to employ the concept of phenomenal consciousness comes from the work of Knobe and Prinz (2008) reviewed in the previous section. Unfortunately, a number of criticisms have been raised concerning their interpretation of their empirical results, raising doubts about this conclusion. Focusing on Knobe and Prinz's second study, the basic difficulty is the same one that we noted above with regard to Huebner (2010): While Knobe and Prinz's results are compatible with the claim that lay people distinguish between mental states that are and that are not phenomenally conscious, they can also be reasonably explained in other ways.

Sytsma and Machery (2009) suggest one obvious alternative, arguing that rather than interpret Knobe and Prinz's two sets of sentences in terms of phenomenality, they can instead be understood in terms of whether or not they ascribe functional roles or behaviors that corporations

¹⁰ Going the other direction, it is worth noting that Huebner found that participants' judgments about pain and happiness for the cyborg cases were less clear, with judgments about pain for these cases being essentially at chance. He argues that these results "sit far less comfortably with the claim that the commonsense understanding of the mind includes a concept of phenomenal consciousness" (Huebner, 2010, 139).

are capable of. A related criticism of Knobe and Prinz's interpretation of their results is given by Adam Arico (2010). He notes that the sets of sentences used by Knobe and Prinz differ in other ways than whether they attribute a phenomenal or a non-phenomenal state to the group agent. And Arico argues that this opens the door for alternative explanations of the results. Most importantly, there are differences in the context provided in Knobe and Prinz's non-phenomenal sentences in comparison to their phenomenal sentences, which is likely relevant to how natural sounding the ascriptions are.¹¹ To illustrate, compare "Acme Corp. intends to release *a new product this January*" with "Acme Corp. is getting depressed." While the italicized portion of the first sentence plausibly provides context that helps explain what is meant by saying that the corporation intends, no such context is provided for the second sentence. This difference can be drawn out with a slight rewrite to the second sentence, such as "Acme Corp. is getting depressed by the negative reviews of the new product."

To better test the finding that lay people are unwilling to attribute phenomenally conscious mental states to group agents, Arico ran a variation on Knobe and Prinz's second study that included three phenomenal sentences and three non-phenomenal sentences with matching context. Following Knobe and Prinz, Arico used locutions like "is feeling" or "is experiencing" to mark sentences as phenomenal. For example, following this heuristic, the first sentence in each of the following pairs is categorized as phenomenal, while the second is categorized as non-phenomenal:

The Housing Authority is feeling angry about the foreclosures.
The Housing Authority is angry about the foreclosures.

McDonalds is feeling upset about the court's recent ruling.
McDonalds is upset about the court's recent ruling.

¹¹ See Sytsma and Machery (2009, Section 3.3) for a discussion of this issue with regard to Knobe and Prinz's fourth study.

Averaging across the sentences, Arico found no significant difference between the phenomenal sentences (mean rating of 5.11 on the scale used by Knobe and Prinz) and the non-phenomenal sentences (mean rating of 5.07) when matching context was provided. Further, the mean rating for each type of sentence was above the neutral point. Arico takes this to cast doubt on Knobe and Prinz's interpretation of their original results: Using sentences modeled on their phenomenal sentences, but with additional context provided, Arico found that people now tended to rate them as sounding natural.

It is unclear what positive conclusion should be drawn from Arico's results, however. The primary reason is that the lack of a significant difference between these sentences might be taken to reflect that lay people do *not* treat locutions like "is feeling" and "is experiencing" as marking a mental state as being phenomenal. If this is correct, then it is hardly surprising that people rated the sentences similarly. This is cold comfort for Knobe and Prinz, however. Suppose that they were to abandon the claim that such locutions mark mental states as being phenomenal. To bring Arico's results into line with their interpretation of their previous results, they would also need to claim that the sentences in each of Arico's pairs were actually ascribing non-phenomenal mental states to the group agents. But this is implausible, as philosophers typically classify mental states like anger as being phenomenally conscious.

A different type of criticism has recently been raised by Mark Phelan, Adam Arico, and Shaun Nichols (2013). They note that Knobe and Prinz are drawing conclusions about how people think about mental states from data concerning how natural sounding they find a range of sentences ascribing mental states to group agents. But, to do so they need to assume that people adopt a *realist* interpretation of these mental state ascriptions: They must assume that "people 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” (707). Phelan and colleagues challenge this assumption, presenting the results of two studies suggesting that people tend to interpret sentences ascribing mental states to group agents in a *distributivist* fashion—people tend to interpret these sentences as ascribing the mental states not to the group itself, but to the people making up the group (or some subset thereof).¹²

In their first study, Phelan and colleagues gave participants one of three sets of four sentences—the first set containing sentences ascribing non-mental states to group agents, the second containing sentences ascribing mental but non-phenomenal states to group agents, and the third containing sentences ascribing phenomenal states to group agents—along with four filler sentences. In the following examples, the first sentence is from the non-mental set, the second from the non-phenomenal set, and the third from the phenomenal set:

After Boeing lost the Army contract, Boeing needed to lay off workers.

After Boeing lost the Army contract, Boeing expected to have to lay off workers.

After Boeing lost the Army contract, Boeing felt anxious about having to lay off workers.

Participants were then asked to replace the underlined instance of the group name with either “it” or “they.” Phelan and colleagues reason that a “they” response would indicate that the person was thinking about the ascription in a distributivist fashion, taking it to apply to the members of the group rather than the group itself. As shown in Figure 3, Phelan and colleagues found that participants were significantly more likely to select “they” than “it” for either type of mental sentences than for the non-mental sentences, but were not significantly more likely to select “they” for the phenomenal sentences than for the non-phenomenal sentences. Phelan and colleagues take this to indicate that lay people tend to understand mental state ascriptions to

¹² For a general discussion of group minds and how people think about them, see Huebner (forthcoming, 2013).

group agents in a distributivist rather than a realist fashion, casting doubt on Knobe and Prinz's interpretation of their results.

This conclusion is further bolstered by the results of Phelan and colleagues' second study. Participants were given one of the three sets of four sentences used in the previous study (non-mental, non-phenomenal, phenomenal). This time, however, rather than ask participants to replace the group name, they were given two alternative versions of each sentence—one giving a distributivist reading and one a realist reading. Participants were then asked to select the alternative sentence that best describes what was meant by the original sentence. For example, for the (phenomenal) sentence “the Secret Task Force felt threatened by the public reaction,” participants were asked to select from the following two alternatives:

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

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

The results are shown alongside those from the first study in Figure 3. Phelan and colleagues found that participants were significantly more likely to select the distributivist alternative than the realist alternative for either type of mental sentences than for the non-mental sentences, but were not significantly more likely to select the distributivist alternative for the phenomenal sentences than for the non-phenomenal sentences. Again, this suggests that lay people tend to interpret mental state ascriptions to group agents in a distributivist fashion—they take them to be attributing the mental states to members of the group, not to the group itself.




Figure 3: Results of studies in Phelan, Arico, and Nichols (2013), showing percentage of distributivist answers broken down by type of sentence.

Recall that Knobe and Prinz predicted that they would find an asymmetry in responses to the two types of sentences in their second study because they think that the corporation is capable of having non-phenomenal mental states but is not capable of having phenomenal mental states. Phelan and colleagues' results suggest that lay people do not tend to interpret sentences ascribing mental states to a group agent in this type of realist fashion: They tend to read such sentences as ascribing the mental states to members of the group, and do so whether the mental states at issue are phenomenal or non-phenomenal. This suggests that something else is driving the asymmetry found by Knobe and Prinz rather than that their participants were applying the phenomenal/non-phenomenal distinction—perhaps one of the alternative explanations offered by Sytsma and Machery (2009) and Arico (2010). Phelan and colleagues offer a related possibility, calling on what they term the *qua members principle*: They argue that the appropriateness of a

mental state ascription to a group agent depends on whether that mental state is appropriately attributed to members of the group *as* members of that group. Phelan and colleagues then argue that Knobe and Prinz's non-phenomenal sentences were more appropriate to members of Acme Corporation *as employees of Acme Corporation* than were their phenomenal sentences.

3. The Negative Case for a Lay Concept of Phenomenal Consciousness

Flipping the tables, Sytsma and Machery (2010) present evidence suggesting a negative answer to our target question: They argue that most lay people do *not* employ the concept of phenomenal consciousness in making mental state attributions. Sytsma and Machery reason that if lay people call on the concept of phenomenal consciousness, then they should tend to treat paradigmatic examples of phenomenally conscious mental states such as seeing red and feeling pain similarly.

To test this, in their first study Sytsma and Machery gave participants a description of one of the two agents shown in Figure 4—a normal human (Timmy) or a simple, non-humanoid robot (Jimmy)—performing one of two tasks. The tasks were designed to elicit the judgment that the human had either seen red or felt pain; the robot was then described as behaving in an analogous manner. In each of the resulting four scenarios, the agent was instructed to manipulate one of three boxes based on color. In two of the four scenarios, that manipulation was successful and the participants were asked whether the agent “saw red,” answering on a 7-point scale anchored at 1 with “clearly no,” at 4 with “not sure,” and at 7 with “clearly yes.” In the other two scenarios, the agent was electrically shocked and participants were asked whether the agent “felt pain,” answering on the same scale. On a second page, participants were asked how they thought other people would answer the question, and on a third page they were asked to explain their answers.



Figure 4: Images of the human (Timmy) and robot (Jimmy) used in the studies in Sytsma and Machery (2010).

Sytsma and Machery found that participants with more than minimal philosophical training, tended to treat the two types of mental states similarly.¹³ As seen in Figure 5, they tended to attribute both seeing red and feeling pain to the human, and to deny both of these states for the robot. In contrast, participants with little to no training in philosophy treated these mental states differently. While they followed the philosophers in tending to attribute both states to the human and in denying that the robot felt pain, they tended to answer that the robot saw red. These findings have subsequently been replicated a number of times, including in the studies reported in Sytsma and Machery (2012b).

¹³ For purposes of this study, “philosophers” were defined as those participants who had some graduate training in philosophy or who had completed, or were in the process of completing, an undergraduate degree with a major in philosophy.




Figure 5: Results of Study 1 in Sytsma and Machery (2010), showing mean responses for philosophers and non-philosophers for each scenario.

On the basis of these results, as well as those of two follow-up studies investigating other mental states (feeling anger, smelling banana), Sytsma and Machery argue that *contra* Knobe and Prinz, lay people do not tend to employ the concept of phenomenal consciousness. The reasoning is that if people were calling on the concept of phenomenal consciousness in classifying mental states, then we would expect them to treat prototypical examples of phenomenally conscious mental states similarly, all else being equal. Thus, insofar as an entity is thought to be phenomenally conscious and shows behavioral cues associated with seeing red and feeling pain, for example, we would expect people to attribute both of those states to the entity; and insofar as the entity is not thought to be phenomenally conscious, we would expect neither

state to be attributed to the entity. But this is not what Sytsma and Machery found for attributions to the robot Jimmy. They found that lay people tended to attribute one state (seeing red) to the robot, but not the other (feeling pain). As Sytsma and Machery conclude, “in clear contrast to philosophers, the folk do not seem to believe that there is something common to all these mental states—namely that they are phenomenal” (2010, 319).

Sytsma and Machery then argued that this finding has significant implications for recent discussion of phenomenal consciousness in the philosophical and scientific literature. While the existence of phenomenally conscious mental states is thought to pose a hard problem, the common justification for the claim that such mental states exist is that they are evident in first-person experiences with states like seeing red and feeling pain. If this is correct, however, then the existence of phenomenally conscious mental states should be obvious to lay people as well, and they should tend to treat such states similarly, all else being equal. But the evidence suggests that they do not.

Assuming for the moment that Sytsma and Machery are correct and that most lay people do not call on the concept of phenomenal consciousness in categorizing mental states, an obvious follow-up question arises: Insofar as people tend to treat some mental states differently (attributing the state of seeing red but not feeling pain to the robot Jimmy, for example), how do they categorize mental states? On the basis of their second and third studies, Sytsma and Machery suggested that lay mental state attributions depend, in part, on whether the mental state at issue is associated with a valence (whether the mental state is thought to have a hedonic value or be either pleasurable or unpleasurable). They then explained the results of their first study, discussed above, in terms of lay people tending to hold that the robot Jimmy is incapable of finding a mental state to be either pleasurable or unpleasurable, and lay people tending to associate feeling pain, but not

seeing red, with a valence. Subsequent results reported in Sytsma (forthcoming) cast doubt on this *valence account*, however. In its place, Sytsma puts forward a *naïve account* of the findings in Sytsma and Machery (2010), arguing that lay people tend to view both colors and pains not as phenomenal qualities, but as mind independent qualities of objects outside of the mind/brain.¹⁴

4. Alternative Explanations

As with the case made by Knobe and Prinz (2008) for an affirmative answer to the question of whether lay people tend to have the concept of phenomenal consciousness, the case for a negative answer presented by Sytsma and Machery (2010) has met with resistance. The most common response has been to argue that the phrase “seeing red” is polysemous, and that in Sytsma and Machery’s study lay people tended to understand it in an *informational sense* (the entity is responsive to the distinction between red things and non-red things) while philosophers tended to understand it in a *phenomenal sense* (the entity has the appropriate phenomenally conscious mental state). This *ambiguity objection* is raised by Huebner (2010, 137), and pushed by Peressini (2013) and Fiala, Arico, and Nichols (2014).

Sytsma and Machery (2010) offered several responses to this objection. Beyond noting that in the absence of further argumentation the criticism is ad hoc, they argued that the distribution of responses they found did not correspond with what one would expect to see if the criticism were correct and that the explanations that participants gave for negative responses did not suggest that they were employing a phenomenal reading. A further reply was given in Sytsma (2009, 2010c, forthcoming). Sytsma argued that the ambiguity objection supposes a theoretical view of colors that lay people typically do not hold. In line with the naïve account

¹⁴ The naïve account builds on the work of Sytsma (2010c) on the folk theory of visual perception, and the work of Reuter (2011) and Reuter, Phillips, and Sytsma (2014) on the folk theory of pain.

noted above, he presented evidence that lay people tend to treat the colors that we are acquainted with in ordinary visual perception as being mind independent qualities of things in the world. But, such a naïve view does not allow for the distinction between the informational and the phenomenal sense of “seeing red.” Thus, insofar as the evidence indicates that lay people tend to hold a naïve view of colors, they are not best interpreted as asserting that the robot saw red in a merely informational sense.

Anthony Peressini (2013) replies to the responses presented in Sytsma and Machery (2010), although not to Sytsma’s subsequent response based on the naïve view of colors. In particular, he argues that the ambiguity objection should not be dismissed as being ad hoc, providing new reason to expect that lay people would diverge from philosophers in adopting an informational reading of “seeing red.” Peressini argues that lay people tend to exclude artifacts like robots from the category of “experiencers,” and suggests that people will be more likely to adopt the informational reading when the phrase is applied to a non-experiencer.

Peressini offers support for the first claim in the first part of a two-part study. Participants were given a description of the categories of experiencers/non-experiencers, then asked to rate each of 20 entities (including a person in a coma, a virus, seaweed, a dolphin, and several types of computer systems, robots, and androids) on whether they are an experiencer, using a seven-point scale anchored at 1 with “clearly not an experiencer,” at 4 with “unsure,” and at 7 with “clearly an experiencer.” The description of the categories of experiences/non-experiencers read as follows:

As we all know, each of us as conscious human beings have an “inner life.” We are aware of things going on around us and inside our minds. In other words, there is something it is like to be each of us at any given moment: the sum total of what we are sensing, thinking, feeling, etc. We are experiencers.

On the other hand, things like thermostats, burglar alarms, and bread machines do not have an inner life: there is not anything it is like to be these objects, despite the fact that they can monitor conditions around them and make appropriate things happen at appropriate times. They are not experiencers.

Peressini found that people tended to judge that the artifacts they were asked about were non-experiencers, including the sophisticated fictional androids C3PO, Data, and R2D2 (mean ratings of 3.21, 3.07, and 2.69 respectively).

The setup of Peressini's study, however, prompts the concern that his materials *taught* participants the new category of "experiencer." Thus, it might be that participants' ratings simply reflect their judgments about the similarity of the entities tested to the examples given—a living human being versus mechanical artifacts (thermostats, burglar alarms, and bread machines). Given these examples, it is perhaps not surprising that participants tended to classify artifacts as "non-experiencers." Further, while Peressini raises doubts about the first study in Sytsma and Machery (2010), he ultimately concludes that they were correct to deny that lay people tend to have the philosophical concept of phenomenal consciousness, at least as I have articulated it above.

Nonetheless, based on the results of the second part of his study, Peressini argues that lay people do "have a concept of some sense of qualia or 'phenomenality'" (22), even if they do not have the philosophical concept I have focused on. In the second part of his study, Peressini gave the same set of participants ten questions related to the philosophical concept of phenomenal consciousness. His analysis focuses on Questions 3 and 6:

3. Can we ever be sure that you see red the way another person does?

6. Imagine that a completely color blind person got an implant that encoded colors in her visual field with numbers indicating colors, so for example, the sky on a clear day was indicated with a number 1 to indicate blue. Would such a person be able to see blue?

Peressini found that participants tended to give negative responses to each of these questions (mean ratings of 3.07 and 3.00 respectively on a 7-point scale). He then argues that one needs to

posit that lay people have something like the concept of phenomenal consciousness to explain these responses. This would seem to be moving too quickly, however, as the responses to these questions can readily be interpreted in other ways. With regard to Question 3, for example, one might take the negative responses to reflect the in-practice difficulty of exactly comparing differences in ability to discern shades of color, perhaps coupled with concerns that “red” is a category term that might reasonably be thought to have subjective boundaries, rather than as indicating that lay people are committed to private and ineffable visual impressions suggestive of phenomenal qualities as Peressini suggests. Further, this alternative explanation is more consistent with the data presented by Sytsma (2010c). And a similar concern can be raised with regard to interpreting the responses to Question 6, where Peressini seems to run together judgments about the occurrence of visual perception with judgments about the nature of visual perception.¹⁵

Brian Fiala, Adam Arico, and Shaun Nichols (2014) offer a variation on the ambiguity objection. They note that the participants in Sytsma and Machery’s original study were not given a way to express that the robot, Jimmy, performed relevant information-processing behavior, except to affirm that Jimmy saw red. Given this, Fiala and colleagues argue that affirmative responses should not necessarily be taken to indicate that participants thought that Jimmy saw red.¹⁶ To test this objection, they ran a variation on Sytsma and Machery’s study in which

¹⁵ To illustrate, consider a scenario similar to that in Question 6 in which a visual scene is described to a blind person, using the word “blue” to describe the color of the sky. Certainly in this case we would not want to say that the blind person was able to see blue (although she was able to hear the word “blue”), but such a judgment would not seem to commit us to any particular view about phenomenal consciousness.

¹⁶ It should be noted that Fiala and colleagues primary concern is not with our target question of whether or not lay people tend to have the concept of phenomenal consciousness. Rather, their critique is given in defense of their *agency model* of mental state attributions (Arico et al., 2011; Fiala, Arico, and Nichols, 2011; Fiala, Arico, and Nichols, 2014). This model posits that a dual-process cognitive system is responsible for mental state attributions (in general, not attributions of one type of mental state or another), with “low-road” processing operating in a fast, automatic, and domain-specific way to categorize entities as AGENTs based on cues like having facial features and displaying interactive behavior to produce a disposition to attribute mental states to that entity. As such, the model

participants selected those descriptions of Jimmy that seemed correct from a set of six, including that Jimmy detected the target color and that Jimmy saw the target color. Fiala and colleagues found that only seven out of 25 participants (28.0%) selected the “saw” option, providing support for their variation of the ambiguity objection. Sytsma (2014) responds by presenting the results of four follow-up studies using the method employed by Fiala and colleagues, but with larger sample-sizes and restricting to non-philosophers. What Sytsma found is that when potential difficulties with the answer choices were corrected, a significant majority of the participants answered that Jimmy saw the target color, despite the fact that they could otherwise indicate that Jimmy performed the color discrimination task by answering that Jimmy identified or located the relevant box.

Another type of objection is raised by Wesley Buckwalter and Mark Phelan (2013), primarily against the valence account put forward in Sytsma and Machery (2010), but also raising difficulties for the evidence they present for the hypothesis that lay people do not tend to have the concept of phenomenal consciousness. While Buckwalter and Phelan target the second and third studies presented by Sytsma and Machery, their concern also applies to Sytsma and Machery’s first study. In line with their contention that lay people tend to be functionalists about mental state attributions (see Phelan and Buckwalter, forthcoming), they argue that asymmetry between the mental state attributions of philosophers and non-philosophers to the robot Jimmy in Sytsma and Machery’s first study might simply reflect that these groups tend to think about the robot differently with regard to its functionality: “That is, perhaps the robot isn’t the same simple

predicts that low-road processing will dispose people to attribute mental states like feeling pain to the robot Jimmy, contrary to what Sytsma and Machery (2010) found. This conflict then motivates their critique of Sytsma and Machery’s experiment. An alternative explanation compatible with the agency model is that Sytsma and Machery’s findings were not driven by low-road processing, but by “high-road” processing—a possibility that is independently suggested by the results reported in Sytsma and Machery (2012b) in response to a distinct criticism put forward by Brian Talbot (2012).

robot for the folk as it is for philosophers.” (360). Phelan and Buckwalter’s suggestion is that non-philosophers might assume that such a robot would be designed to perform functions requiring sensory perception, but not functions requiring the type of bodily monitoring involved in pain. And they hold that this could explain why non-philosophers were willing to attribute seeing red, but not feeling pain, to the simple robot Jimmy.

5. Conclusion

Explaining phenomenal consciousness is considered to be one of the major outstanding problems for modern science. Many hold that it poses a distinctively “hard problem,” resisting our standard forms of scientific explanation. This in turn raises the question of why we should believe in the existence of phenomenal consciousness in the first place. Often the existence of phenomenal consciousness is simply taken for granted, other times it is said to be obvious in ordinary experience. This gives us reason to investigate people’s intuitions about consciousness. While the topic of how people understand and attribute those mental states that philosophers have taken to be phenomenally conscious has not received much attention in the past, that has changed in recent years with a spate of new research in experimental philosophy of mind investigating attributions of consciousness. While the results are mixed, and much research remains to be done, the current evidence suggests that lay intuitions about consciousness are not as obvious as many have assumed.

References

- Alexander, Joshua (2012). *Experimental Philosophy: An Introduction*. Malden, MA: Polity Press.
- Arico, Adam (2010). "Folk psychology, consciousness, and context effects." *Review of Philosophy and Psychology*, 1(3), 371–393.
- Arico, Adam, Brian Fiala, Robert Goldberg, and Shaun Nichols (2011). "The folk psychology of consciousness." *Mind & Language*, 26, 327–352.
- Baron-Cohen, Simon (2001), "Theory of mind in normal development and autism." *Prisme*, 34, 174–183.
- Bering, Jesse (2002). "Intuitive conceptions of dead agents' minds: the natural foundations of afterlife beliefs as phenomenological boundary." *Journal of Cognition and Culture*, 2(4), 263–308.
- Bering, Jesse and David Bjorklund (2004). "The natural emergence of afterlife reasoning as a developmental regularity." *Developmental Psychology*, 40(2), 217–233.
- Bering, Jesse, Carlos Blasi, and David Bjorklund (2005). "The development of 'afterlife' beliefs in religiously and secularly schooled children." *British Journal of Developmental Psychology*, 23(4), 587–607.
- Block, Ned (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), Minneapolis: University of Minnesota Press, pp. 261–326.
- Buckwalter, Wesley and Mark Phelan (2013). "Function and feeling machines: a defense of the philosophical conception of subjective experience." *Philosophical Studies*, 166(2), 349–361.
- Buckwalter, Wesley and Mark Phelan (2014). "Phenomenal consciousness disembodied." In J. Sytsma (ed.), *Advances in Experimental Philosophy of Mind Companion to Experimental Philosophy*, London: Bloomsbury.
- Chalmers, David (1995). "Facing up to the problem of consciousness." *Journal of Consciousness Studies*, 2(3), 200–219.
- Chalmers, David (1996). *The Conscious Mind*. Oxford: Oxford University Press.
- De Brigard, Felipe (2010). "Consciousness, attention and commonsense." *Journal of Consciousness Studies*, 17(9–10), 189–201.
- Dennett, Daniel (1981). *Brainstorms: Philosophical Essays on Mind and Psychology*. Cambridge: MIT Press.
- Dennett, Daniel (1987). *The Intentional Stance*. Cambridge: MIT Press.

Fiala, Brian, Adam Arico, and Shaun Nichols (2011). "On the psychological origins of dualism: dual-process cognition and the explanatory gap." In E. Slingerland and M. Collard (eds.), *Creating Consilience: Issues and Case Studies in the Integration of the Sciences and Humanities*, Oxford: Oxford University Press.

Fiala, Brian, Adam Arico, and Shaun Nichols (2014). "You, robot." In E. Machery and E. O'Neill (eds.), *Current Controversies in Experimental Philosophy*, New York: Routledge.

Fiske, Susan, Amy Cuddy, Peter Glick, and Jun Xu (2002). "A model of (often mixed) stereotype content: competence and warmth respectively follow from status and competition." *Journal of Personality and Social Psychology*, 82(6), 878–902.

Gopnik, Allison and Andrew Meltzoff (1997). *Words, Thoughts, and Theories*. Cambridge: MIT Press.

Gray, Heather, Kurt Gray, and Daniel Wegner (2007). "Dimensions of mind perception." *Science*, 315: 619.

Gray, Kurt, Joshua Knobe, Mark Sheskin, Paul Bloom, and Lis Barrett (2011a). "More than a body: mind perception and the nature of objectification." *Journal of Personality and Social Psychology*, 101(6), 1207–1220.

Gray, Kurt, Adrianna Jenkins, Andrea Heberlein, and Daniel Wegner (2011b). "Distortions of mind perception in psychopathology." *Proceedings of the National Academy of Sciences*, 108(2), 477–479.

Gray, Kurt and Chelsea Schein (2012). "Two minds vs. two philosophies: mind perception defines morality and dissolves the debate between deontology and utilitarianism." *Review of Philosophy and Psychology*, 3, 405–423.

Gray, Kurt, Adam Waytz, and Liane Young (2012a). "The moral dyad: a fundamental template unifying moral judgment." *Psychological Inquiry*, 23(2), 206–215.

Gray, Kurt, Liane Young, and Adam Waytz (2012b). "Mind perception is the essence of morality." *Psychological Inquiry*, 23(2): 101–124.

Gray, Kurt and Daniel Wegner (2009). "Moral typecasting: divergent perceptions of moral agents and moral patients." *Journal of Personality and Social Psychology*, 96(3), 505–520.

Gray, Kurt and Daniel Wegner (2010). "Blaming god for our pain: human suffering and the divine mind." *Personality and Social Psychology Review*, 14(1), 7–16.

Gray, Kurt and Daniel Wegner (2011a). "Morality takes two: dyadic morality and mind perception." In M. Mikulincer and P. Shaver (eds.), *The Social Psychology of Morality: Exploring the Causes of Good and Evil*, Washington, DC: American Psychological Association.

Gray, Kurt and Daniel Wegner (2011b). “To escape blame, don’t be a hero—be a victim.” *Journal of Experimental Social Psychology*, 47(2), 516–519.

Haslam, Nick (2006). “Dehumanization: an integrative review.” *Personality and Social Psychology*, 10(3), 252–264.

Haslam, Nick, Yoshihisa Kashima, Stephen Loughnan, Junqui Shi, and Caterina Suitner (2008). “Subhuman, inhuman, and superhuman: contrasting humans with nonhumans in three cultures.” *Social Cognition*, 26(2), 248–258.

Huebner, Bryce (2010). “Commonsense concepts of phenomenal consciousness: does anyone care about functional zombies?” *Phenomenology and the Cognitive Sciences*, 9(1), 133–155.

Huebner, Bryce (2013). *Macro cognition: A Theory of Distributed Minds and Collective Intentionality*. Oxford: Oxford University Press.

Huebner, Bryce (forthcoming). “The group mind in commonsense psychology.” In J. Sytsma and W. Buckwalter (eds.), *A Companion to Experimental Philosophy*, Oxford: Blackwell.

Huebner, Bryce, Michael Bruno, and Hagop Sarkissian (2010). “What Does the Nation of China Think About Phenomenal States?” *Review of Philosophy and Psychology*, 1(2): 225–243.

Jack, Anthony (forthcoming). “A scientific case for conceptual dualism: the problem of consciousness and the opposing domains hypothesis.” In J. Knobe, T. Lombrozo, and S. Nichols (eds.), *Oxford Studies in Experimental Philosophy (Vol. 1)*, Oxford: Oxford University Press.

Jack, Anthony and Philip Robbins (2012). “The phenomenal stance revisited.” *Review of Philosophy and Psychology*, 3, 383–403.

Jack, Anthony, Philip Robbins, Jared Friedman, and Chris Meyers (2014). “More than a feeling: counterintuitive effects of compassion on moral judgment.” In J. Sytsma (ed.), *Advances in Experimental Philosophy of Mind Companion to Experimental Philosophy*, London: Bloomsbury.

Knobe, Joshua (2008). “Can a robot, an insect or God be aware?” *Scientific American Mind*, 19: 68–71.

Knobe, Joshua and Jesse Prinz (2008). “Intuitions about consciousness: experimental studies.” *Phenomenology and the Cognitive Sciences*, 7: 67–83.

Koch, Christof (2004). *The Quest for Consciousness: A Neurobiological Approach*. Englewood, CO: Roberts & Company.

Machery, Edouard and Justin Sytsma (2011). “Robot pains and corporate feelings.” *The Philosophers’ Magazine*, 1st Quarter: 78–82.

- Malle, Bertram (2005). "Folk theory of mind: conceptual foundations of human social cognition." In R. Hassan, J. Uleman, and J. Bargh (eds.), *The New Unconscious*, New York: Oxford University Press.
- Nagel, Thomas (1974). "What is it like to be a bat?" *The Philosophical Review*, 83: 435–450.
- Onishi, Kristine and Renée Baillargeon (2005). "Do 15-month-old infants understand false beliefs?" *Science*, 308 (5719), 255–258.
- Peressini, Anthony (2013). "Blurring two conceptions of subjective experience: folk versus philosophical phenomenality." *Philosophical Psychology*, DOI: 10.1080/09515089.2013.793150
- Perner, Josef (1991). *Understanding the Representational Mind*. Cambridge: MIT Press.
- Phelan, Mark, Adam Arico, and Shaun Nichols (2013). "Thinking things and feeling things: on an alleged discontinuity in the folk metaphysics of mind." *Phenomenology and the Cognitive Sciences*, 12: 703–725.
- Phelan, Mark and Wesley Buckwalter (forthcoming). "Analytic functionalism and mental state attribution." *Philosophical Topics*.
- Phelan, Mark and Adam Waytz (2012). "The moral cognition/consciousness connection." *Review of Philosophy and Psychology*, 3, 293–301.
- Piazza, Jared, Justin Landy, and Geoffrey Goodwin (forthcoming). "Cruel nature: harmfulness as an important, overlooked dimension in judgments of moral standing." *Cognition*.
- Pinker, Stephen (2007). "The brain: the mystery of consciousness." <http://content.time.com/time/magazine/article/0,9171,1580394,00.html>
- Reuter, Kevin (2011). "Distinguishing the appearance from the reality of pain." *Journal of Consciousness Studies*, 18(9-10): 94–109.
- Reuter, Kevin, Dustin Phillips, and Justin Sytsma (2014). "Hallucinating Pain." In J. Sytsma (ed.), *Advances in Experimental Philosophy of Mind*, London: Bloomsbury.
- Robbins, Philip and Anthony Jack (2006). "The phenomenal stance." *Philosophical Studies*, 127(1), 59–85.
- Stoljar, Daniel (2006). *Ignorance and Imagination: The Epistemic Origin of the Problem of Consciousness*. Oxford: Oxford University Press.
- Sytsma, Justin (2009). "Phenomenological obviousness and the new science of consciousness." *Philosophy of Science*, 76(5), 958–969.

- Sytsma, Justin (2010a). *Phenomenal Consciousness as Scientific Phenomenon? A Critical Investigation of the New Science of Consciousness*. PhD Dissertation, University of Pittsburgh.
- Sytsma, Justin (2010b). “Folk psychology and phenomenal consciousness.” *Philosophy Compass*, 5/8, 700–711.
- Sytsma, Justin (2010c). “Dennett’s theory of the folk theory of consciousness.” *Journal of Consciousness Studies*, 17(3-4), 107–130.
- Sytsma, Justin (2014a). *Advances in Experimental Philosophy of Mind*. London: Bloomsbury.
- Sytsma, Justin (2014b). “The robots of the dawn of experimental philosophy of mind.” In E. Machery and E. O’Neill (eds.), *Current Controversies in Experimental Philosophy*, New York: Routledge.
- Sytsma, Justin (forthcoming). “Revisiting the valence account.” *Philosophical Topics*.
- Sytsma, Justin and Edouard Machery (2009). “How to study folk intuitions about consciousness.” *Philosophical Psychology*, 22(1): 21–35.
- Sytsma, Justin and Edouard Machery (2010). “Two conceptions of subjective experience.” *Philosophical Studies*, 151(2): 299–327.
- Sytsma, Justin and Edouard Machery (2012a). “The two sources of moral standing.” *Review of Philosophy and Psychology*, 3, 303–324.
- Sytsma, Justin and Edouard Machery (2012b). “On the relevance of folk intuitions: a reply to Talbot.” *Consciousness and Cognition*, 21(2): 654–660.
- Sytsma, Justin and Jonathan Livengood (forthcoming). *The New Experimental Philosophy: An Introduction and Guide*. Peterborough, ON: Broadview Press.
- Talbot, Brian (2012). “The irrelevance of folk intuitions to the ‘hard problem’ of consciousness.” *Consciousness and Cognition*, 21(2): 644–650.
- Theriault, Jordan and Liane Young (2014). “Taking an ‘intentional stance’ on moral psychology.” In J. Sytsma (ed.), *Advances in Experimental Philosophy of Mind Companion to Experimental Philosophy*, London: Bloomsbury.
- Waytz, Adam, Kurt Gray, Nicholas Epley, and Daniel Wegner (2010). “Causes and consequences of mind perception.” *Trends in Cognitive Sciences*, 14: 383–388.
- Wellman, Henry (1990). *The Child’s Theory of Mind*. Cambridge: MIT Press.
- Wilkes, Kathleen (1988). “-----, yishi, duh, um and consciousness.” In A. Marcel and E. Bisiach (eds.), *Consciousness in Contemporary Science*, Oxford: Oxford University Press.