


This article was downloaded by: 10.3.97.143

On: 02 Dec 2023

Access details: *subscription number*

Publisher: *Routledge*

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: 5 Howick Place, London SW1P 1WG, UK



The Routledge Handbook Of Consciousness

Rocco J. Gennaro

Consciousness and Experimental Philosophy

Publication details

<https://www.routledgehandbooks.com/doi/10.4324/9781315676982-35>

Chad Gonnerman

Published online on: 26 Mar 2018

How to cite :- Chad Gonnerman. 26 Mar 2018, *Consciousness and Experimental Philosophy from: The Routledge Handbook Of Consciousness* Routledge

Accessed on: 02 Dec 2023

<https://www.routledgehandbooks.com/doi/10.4324/9781315676982-35>

PLEASE SCROLL DOWN FOR DOCUMENT

Full terms and conditions of use: <https://www.routledgehandbooks.com/legal-notices/terms>

This Document PDF may be used for research, teaching and private study purposes. Any substantial or systematic reproductions, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The publisher shall not be liable for an loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

34

CONSCIOUSNESS AND
EXPERIMENTAL PHILOSOPHY*Chad Gonnerman*

Except in our most skeptical moods, it is hard for us philosophers to doubt the successes of science. Reflection on its achievements over the last 50 years alone is cause for optimism (e.g., the sequencing of the human genome, the development of a rich understanding of the architecture of human memory, and the discovery that the universe is rapidly expanding at an ever-increasing rate). It seems that science is quite good at achieving its epistemic ends—its goals with respect to the likes of knowledge, understanding, and explanation. But if we had to settle on one word that best represents the limits of science, at least as currently practiced, we could hardly do better than the one that unites this volume: consciousness.

But not all forms of consciousness are equally intractable. The one that has struck many as particularly resistant to our best theorizing is *phenomenal consciousness*. This is the form exhibited by mental states that have a subjective, qualitative, or experiential aspect to them. There is “something it is like” to be in them (Nagel 1974). At a minimum, these *phenomenal states* include perceptual experiences, bodily sensations, felt emotions, and felt moods (Tye 2016). So, think about the experience of sucking on a lemon, the sensation of stubbing your toe, and the elation of a hard-won A, and you’ll have a pretty good idea of the type of mental state and form of consciousness that has seemed so puzzling to so many.

Just as we may note that some forms of consciousness seem more intractable than others, we may also note that some puzzles about phenomenal consciousness seem more puzzling than others. Certain “how” or “why” questions seem particularly perplexing. Here is how Chalmers (1995: 201) puts the mystery:

The really hard problem is the problem of *experience*. . . Why is it that when our cognitive systems engage in visual and auditory information processing, we have visual or auditory experience: the quality of deep blue, the sensation of middle C? How can we explain why there is something it is like to entertain a mental image or to experience an emotion? It is widely agreed that experience arises from a physical basis, but we have no good explanation of why and how it so arises.

This quote captures what Chalmers calls “the hard problem of consciousness.” It contrasts with what he deems to be the easier problems. These are problems primarily having to do with how the mind/brain accesses and deploys information. For instance, what underwrites our capacity

to discriminate environmental stimuli? The hard problem is the problem of explaining how or why certain mental states feel the way they do—or why they feel any way at all.

When it comes to questions about what science can and cannot reveal about phenomenal consciousness, we are apt to think about the individual undergoing the states. Call him/her the *experiencer*. To illustrate, consider how Churchland responds to the claim that there is a hard problem of consciousness. According to her, we should simply “get on with the task of seeing how far we get when we address neurobiologically the problems of mental phenomena” (1995: 402). Or consider Chalmers’ (2002) argument that there is a hard problem. According to him, neuroscience fails to provide any hints of how or why certain neural happenings are associated with particular experiences. Why do we have an experience of redness, not blueness, when we look at the typical fire engine? Indeed, why is there an experience at all? Why don’t we simply register and process the color as a thermostat registers and processes ambient temperature? What is important for current purposes is that both Churchland and Chalmers are focused on the sciences of the experiencer. They just have different takes on how far the sciences will go.

Goldman and McGrath (2015) outline two ways of using cognitive science in epistemology. There are applications to the epistemic *subject* and applications to the epistemic *attributor*. We can draw a similar distinction here. We might then ask, “What could a science of consciousness *attribution* reveal? What could we learn from a science of the neural and cognitive systems that underwrite our capacities to think about phenomenal states?” Presumably, we could discover how these systems generate our attributions of phenomenal states to ourselves and to others. It would tell us, for instance, about the cognitive mechanisms and processes at play when we judge a person to be in pain. And, on some views of concepts, an account of this sort would also get at the contents, structure(s), and vehicle(s) of our phenomenal concepts. Results like these would be great. But could they also help in trying to understand phenomenal consciousness in general and the relevant phenomenal states in particular? Could they help dominant strands of philosophical research on consciousness? The view of many experimental philosophers is, yes. But, even if they are wrong, as we will see, the experimental philosophy of consciousness is fascinating for many reasons.

This chapter reviews research in this rapidly growing corner of the experimental philosophy of mind (for more on the general area, see Sytsma and Buckwalter 2016). Section 1 explores debates about how to characterize experimental philosophy. It ends by adopting a narrower conception that emphasizes cognitive scientific methods and the intuitions of ordinary people, or “the folk” (i.e., neurally typical adults with no extensive training in philosophy or the cognitive sciences). Section 2 gives an origins story. It discusses two papers that have significantly shaped the experimental philosophy of consciousness. Here, we will see one way in which experimental work can potentially inform dominant strands of philosophical research on consciousness, namely, by buttressing key lines of thought aimed at establishing a hard problem of consciousness. Section 3 gets into experimental research on sentences attributing phenomenal states to entire groups. The research landscape here is very unsettled. Some work suggests that people find these sentences to be acceptable at times, seeming to embrace the existence of group minds with phenomenal states. Other work suggests that people treat these sentences as attributions to individual group members. Finally, Section 4 discusses work looking into folk attributions of phenomenal states to individuals. The emphasis is on accounts of the psychological processes and mechanisms underlying these attributions.

1 A Framework

What is experimental philosophy? While the question has received much attention over recent years, experimental philosophers, evidently, have yet to converge on a single answer. There are

some who give voice to a narrower conception of the area. The claim is that experimental philosophers use the methods of the cognitive sciences to study philosophical intuitions, especially of ordinary people (e.g., Nadelhoffer and Nahmias 2007; Knobe and Nichols 2008; Alexander 2012). There are also some who propose a broader conception. According to these, experimental philosophers use scientific methods, cognitive or not, to study whatever may help advance philosophical research, whether these are intuitions or not (e.g., Rose and Danks 2013; O’Neill and Machery 2014; Sytsma and Livengood 2016). Who is right?

There is some reason to think that the narrower conception is more often pursued. Silber and Knobe surveyed five years of publications in experimental philosophy (Knobe 2016). They culled all the studies and sorted them into two bins. The guiding question was: How did the paper use the study? Did it use the results to make a positive contribution to conceptual analysis? Or did it use the results to make a negative case against the intuitive practices of analytic philosophy? Reflecting on these findings, Knobe (2016) concludes that negative experimental philosophy is pretty rare. Instead, as the title of his paper goes, experimental philosophy is cognitive science. It is in the business of developing and testing hypotheses about the psychology responsible for our philosophical intuitions.

To illustrate the idea, consider Buckwalter and Phelan’s (2014) first study. Participants read a short story about an agent who got in a car wreck while driving to his son’s home in order to leave some photos in the treehouse. In one condition, the agent emerges from the wreck unscathed and walks the photos to the treehouse. In another, the agent dies in the wreck; undeterred, however, he floats the photos to the treehouse as a ghost. Buckwalter and Phelan found that attributions of felt emotions were indistinguishable across the two conditions. Participants were equally willing to attribute felt anger and felt happiness to the agent whether he was a human or a ghost. So, in view of Knobe’s characterization of experimental philosophy, we might expect to see Buckwalter and Phelan give a cognitive–scientific spin to their results. And this is what they do. Buckwalter and Phelan claim that their results provide evidence against the *embodiment hypothesis*, the view “that unified *biological embodiment* is a major psychological factor that cues ordinary attributions of experiences, feelings, emotions, and so on, to other entities” (2014: 46; for a defense of the hypothesis or ones very close to it, see Knobe 2008; Knobe and Prinz 2008; Gray et al. 2011).¹ Thus Buckwalter and Phelan are in line with Knobe’s proposal. They are engaged in what Arico (2010: 372) describes as the “new Folk Psychology project”—that of answering the question “how do we, as humans, go about attributing phenomenal states?”

Knobe’s proposal does a decent job at capturing most work in experimental philosophy. Or at least it does so once we correct for one notable omission. It fails to make it perfectly obvious that experimental philosophy is, well, philosophy. After all, cognitive science is not philosophy. This is not to diminish the role that philosophy plays in cognitive science. And it is not to deny philosophy’s long interest in minds. It is only to acknowledge a common thought: cognitive science is an interdisciplinary area of inquiry that goes beyond the contributing disciplines without amounting to a new discipline (Bechtel 1986). Thus, cognitive science is not philosophy, psychology, linguistics, etc. This claim, combined with Knobe’s proposal as articulated, would imply that experimental philosophy is not philosophy. But this is false. The ‘experimental’ in ‘experimental philosophy’ is not like the ‘fake’ in ‘fake diamonds.’ It’s more like the ‘good’ in ‘good ideas.’ It modifies the noun to identify a subset.

This view is not merely my own. Sytsma and Livengood surveyed philosophers for their views of experimental philosophy. They found nearly 90% agreement that addressing a philosophical problem or issue is needed for a paper to qualify as a work in experimental philosophy (2016: 22). So, we should revise Knobe’s account. A good start would be to say that experimental philosophy is cognitive science that addresses something philosophical. Sometimes, the

two simply overlap. This is especially true if Knobe (2007) is right in claiming that the giving of cognitive accounts represents a return to a traditional conception of philosophy practiced by the likes of Plato, Hume, and Nietzsche. And sometimes the author deploys her cognitive story for (other) philosophical ends.

But, even with this revision, we have yet to nail down an uncontroversial account of experimental philosophy. Many will regard our revised proposal as being too narrow. It fails, they'll insist, to capture all of the activity worthy of the label. This line of thought is especially likely to occur to experimental philosophers. They strongly tend to favor a broad conception of their area (Sytsma and Livengood 2016: 15–18). So, maybe, for the purposes of this chapter, we should adopt a broad conception of the experimental philosophy of consciousness. It would include the scientific study of intuitions about consciousness. It would also include the scientific study of anything else that might reasonably advance philosophical work on consciousness. Some may worry that the result is too broad; we are left an assemblage of work that is not deeply unified. But so what? Many areas of philosophy are loose assemblages, more or less. Still, loose assemblages can make for messy reviews.

The good news is that we can keep this review reasonably tidy, since most of the extant work on consciousness qualifies as experimental philosophy narrowly construed (Sytsma 2010a, 2014a). So, a good chunk of this work involves probing the intuitions of ordinary people to get at the underlying systems, processes, and mechanisms. And, while there are many types of phenomena to which we can apply the word 'consciousness' (Natsoulas 1978; Rosenthal 1986; Block 1995), most experimental philosophical work is on phenomenal consciousness in particular. So, for the purposes of this chapter, I am going to follow the herd, working with the following characterization:

The experimental philosophy of consciousness involves (1) the methodical collection and analysis of empirical data pertaining to ordinary intuitions of, or related to, phenomenal consciousness or phenomenal states (2) using cognitive–scientific methods in order to contribute to research in (3) the cognitive sciences of these intuitions and (4) the philosophy of phenomenal consciousness, phenomenal states, and associated phenomena.

Of course, by working within this narrower framework, I am not denying that there is interesting work outside it. Reuter's (2011) corpus analysis of uses of pain terms, and his claim that people distinguish between pains and pain experiences, is an excellent example of the broader work. Here, I am simply trying to give the reader a decent sense of where the bulk of the literature is at the moment.

2 Origins

One work that has shaped the experimental philosophy of consciousness is Gray, Gray, and Wegner (2007). The paper reports the results of a study in which participants made pairwise comparisons of 13 characters for 18 mental capacities. They had to determine, for instance, "whether a girl of 5 is more or less likely to feel pain than a chimpanzee is" (2007: 619). According to Gray et al., the results indicate that people view minds as differing along two dimensions—Agency and Experience. It seems that from a folk perspective, minds can vary in their capacities for *agential* states (high-level cognitive states like those that figure in planning, communication, and thought) and *experiential* states (phenomenal states like feelings of hunger, pain, and embarrassment). Indeed, if Gray et al. are right, ordinary people recognize three

types of minds: (1) those with high levels of Experience yet low levels of Agency (e.g., a fetus); (2) those that score high in Agency but low in Experience (e.g., God); and (3) those that register high in both (e.g., a man or a woman).

The claim that people distinguish between these three minds is philosophically interesting in many ways. To get at one, consider again the hard problem of consciousness. Chalmers suggests that we cannot respond to the problem by denying the phenomena. The reason is that “Experience is the central and manifest aspect of mental lives,” and this gives experience a “status as an explanandum” (1995: 207). If true, we might expect people to have a (if not the) concept of phenomenal consciousness (Sytsma and Machery 2009, 2010). After all, we have a way of noticing things that are central and manifest in our lives. And such things often get channeled through our concept-formation processes in order to develop semi-stable bodies of information that can aid our thinking about these things down the road. So, the question arises, do people have a concept of phenomenal consciousness? The results of Gray et al. give us some reason to think that they do.

Their findings suggest that people take some beings to score high in Agency and low in Experience. God appears to be an example. This gives us some (as we’ll see, imperfect) reason to think that the folk are willing to embrace the possibility of *purely intentional minds*.² Here, I am thinking about minds that only have intentional states. They never undergo phenomenal ones. Of course, phenomenal states are the stars of this chapter. They include your perceptual experiences, felt bodily states, and felt emotions. Intentional states are mental states that exhibit *intentionality*. They are about things. Examples include your beliefs, desires, and intentions. Notice that if the folk recognize the possibility of a purely intentional mind, then it seems that their attributions of phenomenal states to humans must be sensitive to cues (deemed to be) had by them but not (deemed to be) had by the purely intentional mind. The psychology of phenomenal state attribution must look different than the psychology of intentional state attribution. If so, this is a notable result. Some mindreading research—work on the systems responsible for our mental state attributions to ourselves and to others (for an introduction, see Nichols and Stich 2003)—deliberately base their accounts on findings limited to mindreading vis-à-vis intentional states (e.g., Apperly 2011: 4). Thus they may be only getting at part of the story, while masquerading as authors of the entire story. Moreover, we now have reason to think that ordinary people have a concept of phenomenal consciousness. Positing that they do would go a long way toward explaining why they distinguish between human minds and purely intentional minds. The claim would be that the distinction is driven by a concept that heavily weights cues had by the former and not the latter.

With that said, the underlying argument is not trivial. They rarely are when the goal is to establish a substantive claim about minds. One question we might ask is, “To what extent do the results of Gray et al. show that people view some minds as scoring high in Agency and low in Experience? There are reasons to think that the results are only moderately indicative at best. Notice that the generalization that people view some minds as scoring high in Agency and low in Experience stems from results suggesting that they tend to view God in particular as scoring in these ways. But the set of experiential states explicitly explored by Gray et al. was pretty narrow. It included the likes of hunger, fear, pain, rage, and desire. As Phelan, Arico, and Nichols (2013) note, these aren’t exactly the types of states that seem appropriate for a being like God. How could God feel hungry, for example? It may be that had Gray et al. looked at a broader range of states, including love and maybe even anger, God would have come out looking rather differently (for a related worry, see Sytsma and Machery 2010). This line of thought raises the question of whether there is any additional evidence that people have a concept of phenomenal consciousness. Here, work by Knobe and Prinz might help.

Knobe and Prinz (2008) report the results of five studies. In one of the most discussed, they gave participants four sentences:

- a Acme Corp. is feeling upset.
- b Acme Corp. is upset about the court's recent ruling.
- c Acme Corp. is feeling regret.
- d Acme Corp. regrets its recent decision.

In (a) and (c), 'feeling' is used to attribute a mental state to a group agent. In (b) and (d), this word is not used, but a similar, if not identical, state is attributed to the group. Participants were asked to assess each sentence according to how weird or natural it sounded. They tended to say that (a) and (c) sounded weird while (b) and (d) sounded normal. Knobe and Prinz contend this shows that ordinary people have a concept of phenomenal consciousness. Their thought seems to go: that the results pattern as they do suggests that people are unwilling to attribute phenomenal states to group agents while perfectly willing to attribute intentional (or at least non-phenomenal) states to these agents. So, there is evidence that, from a folk perspective, group minds are purely intentional minds. We have thus landed on a discovery that calls for explanation. Knobe and Prinz maintain that the best explanation involves positing a folk concept. They suggest that when trying to determine whether a target is in phenomenal state P (say, a state of felt regret), people pull two concepts from long-term memory (LTM). The first carries information about the typical causes and effects of P—its *functional profile* (e.g., for regret, it might include increased rumination, feelings of self-blame, and a sense of loss for what could have been). The second concept pulled from LTM is one for phenomenal consciousness. Importantly, it is a concept that requires the target to have a spatially unified, physical body. Thus, according to Knobe and Prinz, the reason a sentence like 'Acme Corp. is feeling regret' rings odd to us is that, in building an interpretation of the sentence, we are forced to apply our concept of phenomenal consciousness to a target—here, a group agent—that we represent as not meeting the conditions of this concept.

But maybe this argument goes too fast. To underscore one worry, notice that the results reported by Knobe and Prinz are mainly about patterns of *linguistic* intuitions. Again, participants assessed sentences for how weird or natural they sounded. Thus, the extent to which folk concepts drive their results is unclear. It may be that the linguistic assessments stem, at least in part, from rather low-level syntactic processing, including the ease or difficulty that participants had in parsing the sentence. There is evidence that syntactic processing can be very sensitive. This includes attuning to *subcategory preference* information (i.e., the perceived likelihood that a given verb and parsing go hand-in-hand; see Garnsey et al. 1997). The possibility that Knobe and Prinz's data partly stem from syntactic processing is particularly relevant since the sentences they used were not *minimal pairs*. The 'feeling' sentences (a) and (c) above lacked "contextualizing" information had by the non-'feeling' sentences (b) and (d)—information that helps to specify what is being attributed. This detail matters. Arico (2010) reports that participants were more likely to say that a sentence ascribing to a group agent a mental state with contextualizing information sounds more natural than a similar sentence without this information, and this pattern emerged whether the 'feeling' locution was used or not (see also Sytsma and Machery 2009). It may be that what we are seeing here is, in part, an attunement to category frequency information of the sort explored in language sciences.

The early results don't quite establish a folk concept of phenomenal consciousness. Thus, as far as these results go, the status of the apparent upshot of a Chalmers-style argument for a hard problem of consciousness—namely, that there should be some such concept—remains unclear.

Still, the early results did help to launch a wide range of work in the experimental philosophy of consciousness. In the next two sections, I review work in two strands of this subsequent research—one on the folk psychology of group minds and the other on phenomenal state attributions in general. Of course, there are other strands of research. They include (1) work exploring potential connections between phenomenal state ascriptions and perceptions of moral patiency (for overview, see Theriault and Young 2014); (2) research probing ordinary conceptions of particular states commonly taken to be phenomenal such as pain (Sytsma 2010b; Reuter, Phillips, and Sytsma 2014; Sytsma and Reuter forthcoming); and (3) investigations looking into the role that consciousness plays in the folk concepts of free will and moral responsibility (e.g., Shepherd, 2012, 2015). Regrettably, space limitations prevent me from getting into these research strands.

3 Group Phenomenality

The first strand of research that the early results helped to launch is a contribution to the folk psychology of group minds (for review, see Huebner 2014: ch. 5). Research into this part of ordinary psychology is not new. Bloom and Veres (1999), for instance, found that people readily use intentional terms to describe the structured movements of groups as simple as three dots. What helps to set the philosophical work apart is its attention to intuitions of phenomenal states. Huebner, Bruno, and Sarkissian (2010), for example, report that both American and Chinese students found sentences ascribing phenomenal states to groups to sound less natural than sentences ascribing similar states to individuals; the difference, however, was less prominent with their Chinese participants. Huebner et al. go on to suggest that cultural factors may moderate a willingness to ascribe phenomenal states to groups. If so, we would have reason to be wary of Block's (1978) Chinese Nation thought experiment, which asks us to imagine the entire population of China duplicating a person's functional profile. We might worry that the intuitions it elicits are sensitive to cultural backgrounds in ways that are inappropriate for their philosophical deployment in arguments against functionalism (for discussion, see Nado 2014). Arico (2010) reports a similar pattern. He found that people were less willing to say that sentences ascribing phenomenal states to groups sounded natural than sentences ascribing these states to individuals. Yet, unlike in Huebner et al.'s results, we see clear a tendency in Arico's to treat phenomenal attributions to groups as natural sounding, at least when they include contextualizing information (e.g., 'McDonalds is feeling upset about the court's recent ruling' vs. 'McDonalds is feeling upset').

Results like these raise the question of how people interpret sentences ascribing phenomenal states to groups. Do they construe them literally? If not, what is the nonliteral reading? Arico et al. (2011) present the results of a pilot study that begins to examine these questions. After cutting participants who seemed to struggle with the literal-figurative distinction, they report that there was a decreased willingness to rate sentences attributing phenomenal states to groups as "literally true" compared to sentences attributing intentional states to groups. It appears that people are less willing to interpret phenomenal attributions to groups literally than intentional ones. Yet this doesn't quite answer our questions. A decrease in willingness is not a flat-out unwillingness. I may be less eager to eat a cookie after I just had one, but if history is any guide, I'm still game for another one. So, Arico et al.'s report is consistent with an overall, weakened tendency to interpret group phenomenal state ascriptions literally. We'll have to look elsewhere for more probative results.

Phelan et al. (2013) provide strong evidence on this score. For one study, they designed a pronoun replacement task. Participants saw a series of sentences with a dependent clause followed by an independent clause. Each clause had the same subject, ostensibly referring to a group agent. The task was to decide whether the subject in the independent clause is best replaced by a singular or

plural pronoun. To illustrate, does ‘it’ or ‘they’ work best as a replacement for the underlined phrase: ‘When MADD’s Drunk Driving Prevention Act failed, MADD got extremely depressed?’ Phelan et al. found that participants were more likely to pick ‘it’ for nonmental state ascriptions than for intentional and phenomenal state ascriptions. They take this as a sign that people tend to interpret group mental state ascriptions nonliterally, treating them *distributively*. That is, to a first approximation, their claim is that people usually walk away with a reading of such sentences in which the mental state is attributed to *each* group member, not the group itself.

Phelan et al.’s distributive view is certainly reasonable. But it also has some substantial empirical commitments. The traditional picture is that we settle on nonliteral readings only after our initial attempt at interpretation fails to deliver a literal reading that makes sense in the context (e.g., Clark and Lucy 1975). It is arguable that, in experimental contexts like those in Phelan et al.’s study, a literal, collective reading of group mental state ascriptions would make sense. Since their view is that people tend to walk away with a nonliteral reading, they are arguably committed to denying the traditional picture. But maybe this isn’t a problem. There are many reasons to reject the traditional picture (e.g., Gibbs 1983). What may be more problematic is that Phelan et al. seem committed to predicting that we will see no differences in ordinary assessments of group mental state ascriptions and the corresponding distributive sentences. After all, if they are right, people tend to interpret the former as the latter. So, their respective assessments should match, modulo experimental noise. Do they? Not always, according to Jenkins et al. (2014). In response to their group-only vignettes, participants tended to agree that the group had the mental state at issue; however, they tended to disagree when asked whether any or each of the group members had the state (see also Waytz and Young 2012).

Stepping back and looking at the empirical record, I would wager that the psychology of how people interpret group mental state ascriptions is as complicated and variable as any other part of the cognitive science of language interpretation. Bringing this more general work to bear on group mental state ascriptions is likely to reveal similar principles. Yet there may be interesting differences as well. This is one issue that remains wide open in the experimental philosophy of consciousness.

4 Individual Phenomenality

The second strand of research that I want to highlight is work responding to Sytsma and Machery (2010). In this paper, the authors argue against Chalmers’ approach to motivating the hard problem of consciousness. They contend that if the approach is correct, there should be a folk concept of phenomenal consciousness, but there is no such concept. To get to their negative conclusion, Sytsma and Machery asked participants to consider a simple robot, Jimmy. As part of a psychological experiment, Jimmy was put in a room that contained three boxes—red, blue, and green. In one condition, Jimmy correctly executed an order to put the red box in front of the door. In another, Jimmy received an electric shock as he grasped the box, after which he dropped it and backed away. While participants tended to affirm that Jimmy saw red, most denied that he felt pain. Sytsma and Machery take this as evidence that there is no folk concept of phenomenal consciousness. As articulated in Sytsma (2014a), the idea seems to be that if there is a folk concept, people should deploy it when trying to determine whether an entity is in a phenomenal state. And so if people tend to withhold attributions of pain to an entity that displays behaviors associated with this state, they should also tend to withhold attributions of seeing red even if the entity displays behaviors associated with this state. Yet this is not what we see. Again, the overall tendencies are to say that Jimmy saw red but did not feel pain (for similar results, see Sytsma and Machery 2012). Therefore, there is no folk concept. It seems

that Chalmers' approach to establishing that there is a hard problem of consciousness needs some rethinking.

If ordinary people don't have a concept of phenomenal consciousness, then how do they go about attributing phenomenal states? The positive thesis of Sytsma and Machery (2010) is a response to this question. According to them, people attend to considerations of valence. If the mental state is associated with a valence (with being either pleasurable or unpleasurable), people will tend to attribute the state only if they hold that the target is capable of valenced states (with finding a mental state to be pleasurable or unpleasurable); otherwise, people will tend to attribute the state so long as the target shows behaviors associated with the state. Feelings of pain are associated with a valence. And so, Sytsma and Machery propose, the reason people tend to say that Jimmy did not feel pain is that they view simple robots as incapable of finding mental states to be pleasurable or unpleasurable. Matters are different with seeing red. While there may be situations in which we find experiences of seeing red to be pleasant or unpleasant (e.g., in a painting), we don't usually do so, at least not often enough for an association to form. And so, people tend to say that a simple robot like Jimmy sees red as long as it displays behaviors of the right sort.

Not everyone agrees with Sytsma and Machery. Buckwalter and Phelan (2013) respond to the positive thesis. They propose that ordinary attributions of mental states to a simple robot are sensitive to design details. Maybe when people consider the Jimmy vignettes, they ask themselves, "Has the robot been designed to carry out tasks associated with seeing red or feeling pain?" Buckwalter and Phelan present evidence that this question does matter. In other works (Phelan and Buckwalter 2012; Buckwalter and Phelan 2014), they argue that attributions of mental states are sensitive to functional details in a different sense as well. It appears that information about perceptual inputs, behavioral outputs, and other mental states are important to folk attributions, including phenomenal ones (cf. Huebner 2010). Thus, to the extent that Jimmy is thought to be incapable, say, of a wide range of behaviors associated with being in pain, we should see a tendency to withhold attributions of pain to the robot (for a similar suggestion, see Talbot 2012; cf. Sytsma and Machery 2012). And the results of Buckwalter and Phelan support this line of thought.

Another line of criticism comes from Fiala, Arico, and Nichols (2014). In other works, they defend a general picture of the psychology of ordinary mental state attributions (Arico et al. 2011; Fiala, Arico and Nichols 2012). It is a picture that portrays the underlying processes as responsive to rather low-level features. If the target has face-like qualities, displays interactive behavior, or moves in non-inertial ways, then people will be disposed to attribute a wide range of phenomenal and intentional states to the target. This is the Agency Model. According to Fiala et al., the model predicts that people will tend to attribute states of seeing red and feeling pain to Jimmy. But, again, they don't. Why is that? Fiala et al. suggest that there is another stream of processing at play. They characterize it as a "high road" pathway. It involves slow, deliberate, and introspectively accessible reasoning. And, according to them, it produces the widespread response that Jimmy does not feel pain. Their thought is that the high-road pathway overrides the deliverances of Agency-based processing. As they put it, "It is effectively a platitude in our culture that robots are incapable of pain or emotion" (2014: 37). But then why don't we see similar results with Jimmy seeing red? According to Fiala et al., what keeps participants from registering a denial here is the forced-choice nature of the response options used by Sytsma and Machery. Participants sense that Jimmy *detects* red, but the best that these options allow for is to affirm that Jimmy *sees* red. Fiala et al. report that when you give participants the option to say that Jimmy detected the color, they are unlikely to say that Jimmy saw it. But this result may itself be a product of the response options according to Sytsma (2014b). Once these issues are fixed, he finds that people are, once again, generally willing to say that Jimmy saw the color.

Still, Fiala et al. (2014) helps to bring out an important issue. A lot of work in the experimental philosophy of consciousness is driven by responses to vignettes involving simple robots. We might wonder about the extent to which these responses reflect processing of the sort discussed in this section. When participants read a story about a simple robot, they probably build a representation of the target that captures various details given in the story. But it probably includes other properties as well, such as ones deemed to be typical for simple robots. Consider Jimmy again. Bets are that you pictured him made of metal, not cardboard. And I'd wager that a CPU comes to mind before beer cans. Importantly, the properties automatically imputed to Jimmy could include an absence of mentality to some degree or other. If so, then responses to questions about his mental states may be affected by processes that extract information from participant-enriched representations of the target. To the extent that they are, cognitive accounts that depict the underlying processes more generally, as operating regardless of the type of entity in question, are apt to go awry. They'll depict participant responses as telling us, as it were, about how people apply psychological predicates—that they attune to details of valence, function, Agency, etc.—when in fact enriched representations of the subject are doing the heavy lifting. So, what we have here is a potential source of noise that has not received much attention by experimental philosophers of consciousness (for a related discussion of stereotypes in mindreading, see Epley 2014). Hopefully, it will receive more attention down the road.

While there is other work critical of the positive thesis of Sytsma and Machery, including Sytsma's own criticisms (2012), I want to briefly highlight some of the pushback given to their negative thesis. Is there a folk concept of phenomenal consciousness? Talbot (2012) notes that many approaches to this question focus on folk intuitions about the mental states of others, including Sytsma and Machery's approach. In his view, such approaches are ill suited for this end. Peressini (2014) adds to the exchange by arguing that there is a folk concept of phenomenal consciousness. But instead of probing intuitions elicited by short stories, he emphasizes more general intuitions. As interesting as these lines of thought are, in the space remaining, I want to raise one further possibility: perhaps the negative thesis is not all that surprising or troubling.

If we are looking for a concept whose content hews very closely to the philosopher's, maybe we shouldn't be too surprised to find that there is no folk concept of phenomenal consciousness. Again, the philosophical concept is of a mental state for which there is something it is like to be in the state. It covers the likes of biting into a lemon, stubbing a toe, and feeling regret. What it is like being in any of these states is different than being in any of the others. Still, they have something in common. They share the second-order property of there being something it's like to be in them. Should we expect ordinary people to have a concept that explicitly recognizes this commonality? If someone is remarkably introspective, sure, she might note it. But most people? In what arena of ordinary life is it important to draw a systematic distinction between mental states for which there is something it's like to be in them and for which there isn't? I would be surprised if there were a folk concept of this sort. And I am not so sure that it'd be much trouble for the hard problem of consciousness if there weren't. At first blush, it is a problem about phenomenal states, not their concept(s). If there is no folk concept, maybe phenomenal characters are not as central and manifest as some claim. But they don't seem to be entirely foreign to people either. When a doctor asks her patient whether her back pain is dull or sharp, few struggle to understand the question. And few are surprised to discover that pains can feel different ways. None of this is to say that experimental philosophical work on consciousness is uninteresting. Far from it! And it is not to say that this work cannot shed light on the hard problem. Fiala et al. (2012) help to show that it may by arguing that our sense that there is something hard about consciousness stems from a quirk

in our psychology—materials of the sort that figure in physicalist explanations fail to trigger low-level mindreading processes that could otherwise intuitively corroborate the explanation, leaving us with a feeling that the explanation misses something. Instead, all that I am saying is that the connection between the experimental work and the hard problem is less direct than suggested by Sytsma and Machery.

5 Conclusion

We could skim only the surface of the experimental philosophy of consciousness in this chapter. Not only couldn't we give the works discussed herein their full due; there are many that we couldn't discuss. Among these are works more critical of experimental philosophy. The most forceful are works pushing the *expertise defense*—roughly put, the claim that folk intuitions about philosophical matters are less reliable than those of professional philosophers (for review, see Alexander 2016; see also Weinberg et al. 2010). The young experimental philosopher may be heartened to hear that often these criticisms turn on empirical claims that may be experimentally explored and supported—or not, as in some explorations of the expertise defense (e.g., Schwitzgebel and Cushman 2012). What this helps to suggest is that there is still a lot of exciting work to be done in experimental philosophy, including the experimental philosophy of consciousness.³

Notes

- 1 It is unclear to what extent these results falsify the hypothesis. Applying ideas that Cornwell, Barbey, and Simmons (2004) develop for other ends, ghosts are often depicted as affecting the world. We see this in Buckwalter and Phelan's story. This may trigger a sense that the ghost is embodied to some extent, in some way. Indeed, if our higher cognitive capacities are grounded in perceptual-motor simulations, as advocates of embodied cognition argue (Wilson and Foglia 2017), it may be psychologically impossible for us to fully and completely represent a person-like entity as disembodied.
- 2 This inference presupposes that the folk (implicitly) reject a version of the *phenomenal intentional theory* according to which all *possible* intentional states are "constituted by" phenomenal states. Also, when it comes to phenomenally tinged intentional states (e.g., an ordinary experience of seeing a bright red apple or maybe a thought associated with a feeling of understanding), the inference presupposes that the folk view these phenomenal characters as inessential to the intentional states, or at least they must be willing to say that a purely intentional being has intentional states even if its states are deeply unlike ours. For more on phenomenal intentional theory and cognitive phenomenology, see Bourget and Mendelovici (2017).
- 3 The author wishes to thank Morgan Dale, Jacob Robbins, and especially Rocco Gennaro for their helpful comments on an earlier version of this chapter.

References

- Alexander, J. (2012) *Experimental Philosophy: An Introduction*, Cambridge: Polity Press.
- Alexander, J. (2016) "Philosophical Expertise," in J. Sytsma and W. Buckwalter (eds.) *A Companion to Experimental Philosophy*, Malden, MA: Wiley.
- Apperly, I. (2011) *Mindreaders: The Cognitive Basis of "Theory of Mind"*, New York: Psychology Press.
- Arico, A. (2010) "Folk Psychology, Consciousness, and Context Effects," *Review of Philosophy and Psychology* 1: 371–393.
- Arico, A., Fiala, B., Goldberg, R., and Nichols, S. (2011) "The Folk Psychology of Consciousness," *Mind and Language* 26: 327–352.
- Bechtel, W. (1986) "The Nature of Scientific Integration," in W. Bechtel (ed.) *Integrating Scientific Disciplines*, Dordrecht: Martinus Nijhoff.
- Block, N. (1978) "Troubles with Functionalism," *Minnesota Studies in the Philosophy of Science* 9: 261–325.

- Block, N. (1995) "On a Confusion about the Function of Consciousness," *Behavioral and Brain Sciences* 18: 227–247.
- Bloom, P., and Veres, C. (1999) "The Perceived Intentionality of Groups," *Cognition* 71: B1–B9.
- Bourget, D., and A. Mendelovici (2017) "Phenomenal Intentionality," in E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition). <https://plato.stanford.edu/archives/spr2017/entries/phenomenal-intentionality>.
- Buckwalter, W., and Phelan, M. (2013) "Function and Feeling Machines: A Defense of the Philosophical Conception of Subjective Experience," *Philosophical Studies* 166: 349–361.
- Buckwalter, W., and Phelan, M. (2014) "Phenomenal Consciousness Disembodied," in J. Sytsma (ed.) *Advances in Experimental Philosophy of Mind*, London: Bloomsbury.
- Chalmers, D. (1995) "Facing Up to the Problem of Consciousness," *Journal of Consciousness Studies* 2: 200–219.
- Chalmers, D. (2002) "The Puzzle of Consciousness Experience," *Scientific American* 12: 90–99.
- Churchland, P. S. (1995) "The Hornswoggle Problem," *Journal of Consciousness Studies* 3: 402–408.
- Clark, H. H., and Lucy, P. (1975) "Understanding What Is Meant from What Is Said: A Study in Conversationally Conveyed Requests," *Journal of Verbal Learning and Verbal Behavior* 14: 56–72.
- Cornwell, B. R., Barbey, A. K., and Simmons, W. K. (2004) "The Embodied Bases of Supernatural Concepts," *Behavioral and Brain Sciences* 27: 735–736.
- Epley, N. (2014) *Mindwise: How We Understand What Others Think, Believe, Feel, and Want*, New York: Alfred A. Knopf.
- Fiala, B., Arico, A., and Nichols, S. (2012) "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, B., Arico, A., and Nichols, S. (2014) "You, Robot," in E. Machery and E. O'Neill (eds.) *Current Controversies in Experimental Philosophy*, New York: Routledge.
- Garnsey, S. M., Pearlmutter, N. J., Myers, E., and Lotocky, M. A. (1997) "The Contributions of Verb Bias and Plausibility to the Comprehension of Temporarily Ambiguous Sentences," *Journal of Memory and Language* 37: 58–93.
- Gibbs, R. W. (1983) "Do People Always Process Literal Meanings of Indirect Requests?" *Journal of Experimental Psychology: Learning, Memory, and Cognition* 9: 524–533.
- Goldman, A. I., and McGrath, M. (2015) *Epistemology: A Contemporary Introduction*, New York: Oxford University Press.
- Gray, H. M., Gray, K., and Wegner, D. (2007) "Dimensions of Mind Perception," *Science* 315: 619.
- Gray, H. M., Knobe, J., Sheskin, M., Bloom, P., and Barrett, L. B. (2011) "More than a Body: Mind Perception and the Nature of Objectification," *Journal of Personality and Social Psychology* 101: 1207–1220.
- Huebner, B. (2010) "Commonsense Concepts of Phenomenal Consciousness: Does Anyone Care about Functional Zombies?" *Phenomenology and the Cognitive Sciences* 9: 133–155.
- Huebner, B. (2014) *Macrocognition: A Theory of Distributed Minds and Collective Intentionality*, New York: Oxford University Press.
- Huebner, B., Bruno, M., and Sarkissian, H. (2010) "What Does the Nation of China Think about Phenomenal States?" *Review of Philosophy and Psychology* 1: 225–243.
- Jenkins, A. C., Dodell-Feder, D., Saxe, R., and Knobe, J. (2014) "The Neural Bases of Directed and Spontaneous Mental State Attributions to Group Agents," *PLoS ONE* 9(8): e105341. doi:10.1371/journal.pone.0105341
- Knobe, J. (2007) "Experimental Philosophy and Philosophical Significance," *Philosophical Explorations* 10: 119–121.
- Knobe, J. (2008) "Can a Robot, an Insect or God Be Aware?" *Scientific American Mind* 19: 68–71.
- Knobe, J. (2016) "Experimental Philosophy Is Cognitive Science," in J. Sytsma and W. Buckwalter (eds.) *A Companion to Experimental Philosophy*, Malden, MA: Wiley.
- Knobe, J., and Nichols, S. (2008) "An Experimental Philosophy Manifesto," in J. Knobe and S. Nichols (eds.) *Experimental Philosophy*, Oxford: Oxford University Press.
- Knobe, J., and Prinz, J. (2008) "Intuitions about Consciousness: Experimental Studies," *Phenomenology and Cognitive Science* 7: 67–83.
- Nadelhoffer, T., and Nahmias, E. (2007) "The Past and Future of Experimental Philosophy," *Philosophical Explorations* 12: 123–149.
- Nado, J. (2014) "The Role of Intuition," in J. Sytsma (ed.) *Advances in Experimental Philosophy of Mind*, London: Bloomsbury.

- Nagel, T. (1974) "What Is It Like to Be a Bat?" *Philosophical Review* 83: 435–450.
- Natsoulas, T. (1978) "Consciousness," *American Psychologist* 33: 906–914.
- Nichols, S., and Stich, S. (2003) *Mindreading: An Integrated Account of Pretense, Self-Awareness, and Understanding Other Minds*, Oxford: Oxford University Press.
- O'Neill, E., and Machery, E. (2014) "Experimental Philosophy: What Is It Good For?" in E. Machery and E. O'Neill (eds.) *Current Controversies in Experimental Philosophy*, New York: Routledge.
- Peressini, A. (2014) "Blurring Two Conceptions of Subjective Experience: Folk versus Philosophical Phenomenality," *Philosophical Psychology* 27: 862–889.
- Phelan, M., Arico, A., and Nichols, S. (2013) "Thinking Things and Feeling Things: On an Alleged Discontinuity in Folk Metaphysics of Mind," *Phenomenology and the Cognitive Sciences* 12: 703–725.
- Phelan, M., and Buckwalter, W. (2012) "Analytic Functionalism and Mental State Attributions," *Philosophical Topics* 40: 129–154.
- Reuter, K. (2011) "Distinguishing the Appearance from the Reality of Pain," *Journal of Consciousness Studies* 18: 94–109.
- Reuter, K., Phillips, D., and Sytsma, J. (2014) "Hallucinating Pain," in J. Sytsma (ed.) *Advances in Experimental Philosophy of Mind*, London: Bloomsbury.
- Rose, D., and Danks, D. (2013) "In Defense of a Broad Conception of Experimental Philosophy," *Metaphilosophy* 44: 512–532.
- Rosenthal, D.M. (1986) "Two Concepts of Consciousness," *Philosophical Studies* 49: 329–359.
- Schwitzgebel, E., and Cushman, F. (2012) "Expertise in Moral Reasoning? Order Effects on Moral Judgment in Professional Philosophers and Non-Philosophers," *Mind and Language* 27: 135–153.
- Shepherd, J. (2012) "Free Will and Consciousness: Experimental Studies," *Consciousness and Cognition* 21: 915–927.
- Shepherd, J. (2015) "Consciousness, Free Will, and Moral Responsibility: Taking the Folk Seriously," *Philosophical Psychology* 28: 929–946.
- Sytsma, J. (2010a) "Folk Psychology and Phenomenal Consciousness," *Philosophy Compass* 5: 700–711.
- Sytsma, J. (2010b) "Dennett's Theory of the Folk Theory of Consciousness," *Journal of Consciousness Studies* 17: 107–130.
- Sytsma, J. (2012) "Revisiting the Valence Account," *Philosophical Topics* 40: 179–198.
- Sytsma, J. (2014a) "Attributions of Consciousness," *WIREs Cognitive Science* 5: 635–648.
- Sytsma, J. (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, J., and Buckwalter, W. (eds.) (2016) *A Companion to Experimental Philosophy*, Malden, MA: Wiley.
- Sytsma, J., and Livengood, J. (2016) *The Theory and Practice of Experimental Philosophy*, Peterborough, Ontario: Broadview.
- Sytsma, J., and Machery, E. (2009) "How to Study Folk Intuitions about Phenomenal Consciousness," *Philosophical Psychology* 22: 21–35.
- Sytsma, J., and Machery, E. (2010) "Two Conceptions of Subjective Experience," *Philosophical Studies* 151: 299–327.
- Sytsma, J., and Machery, E. (2012) "On the Relevance of Folk Intuitions: A Commentary on Talbot," *Consciousness and Cognition* 21: 654–660.
- Sytsma, J., and Reuter, K. (forthcoming) "Experimental Philosophy of Pain," *Journal of Indian Council of Philosophical Research*.
- Talbot, B. (2012) "The Irrelevance of Folk Intuitions to the 'Hard Problem' of Consciousness," *Consciousness and Cognition* 21: 644–650.
- Theriault, J., and Young, L. (2014) "Taking an 'Intentional Stance' on Moral Psychology," in J. Sytsma (ed.) *Advances in Experimental Philosophy of Mind*, London: Bloomsbury.
- Tye, M. (2016) "Qualia," in E.N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition). <https://plato.stanford.edu/archives/win2016/entries/qualia/>.
- Waytz, A., and Young, L. (2012) "The Group-Member Mind Trade-Off: Attributing Minds to Group versus Group Members," *Psychological Science* 23: 77–85.
- Weinberg, J. M., Gonnerman, C. Buckner, C., and Alexander, J. (2010) "Are Philosophers Expert Intuiters?" *Philosophical Psychology* 23: 33–55.
- Wilson, R. A., and Foglia, L. (2017) "Embodied Cognition," in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition). <https://plato.stanford.edu/archives/spr2017/entries/embodied-cognition/>.

Related Topics

Materialism
Dualism
Consciousness and Intentionality
Robot Consciousness
Consciousness, Free Will, and Moral Responsibility
Consciousness and Emotion