Consciousness, Phenomenal Consciousness, and Free Will

Justin Sytsma and Melissa Snater[1]

Abstract: One body of research in experimental philosophy indicates that non-philosophers by and large do not employ the concept of phenomenal consciousness. Another body of research, however, suggests that people treat phenomenal consciousness as essential for having free will. In this chapter, we explore the tension between these findings. We suggest that the dominant, ordinary usages of 'consciousness' concern notions of being awake, aware, and exercising control, all of which bear a clear connection to free will. Based on this, we argue that findings purporting to show that people take the capacity for phenomenal consciousness to be necessary for free will are better interpreted in terms of a non-phenomenal understanding of consciousness. We explore this suggestion by calling on extant work on the dimensions of mind perception, and we expand on it, presenting the results of a new study employing a global sample.

There is a tension between some recent work in experimental philosophy concerning how non-philosophers conceptualize mental states and work on the relationship between attributions of consciousness and free will. While researchers are not in complete agreement about how 'phenomenal consciousness' should be understood, the standard idea is that a mental state is phenomenally conscious just in case there is "something it is like" (Nagel 1974) to be in that state, with this being understood in terms of the state having phenomenal qualities. A growing body of evidence, however, suggests that non-philosophers do not tend to employ this philosophical concept (e.g., Sytsma and Machery 2010, Sytsma and Ozdemir 2019). This work shows that non-philosophers do not tend to categorize mental states in the way that philosophers do, distinguishing between those states that are phenomenally conscious and those that are not. At the same time, another body of evidence has been taken to indicate that non-philosophers typically treat the possession of phenomenal consciousness as a necessary condition for freely willed action (e.g., Shepherd 2015, Nahmias et al. 2020). But if non-philosophers don't employ the concept of phenomenal consciousness, then they couldn't treat phenomenal consciousness as a necessary condition for free will; and if they do treat it as a necessary condition, then researchers arguing they don't employ such a concept must be mistaken.

[1] Forthcoming in *Advances in Experimental Philosophy of Action* edited by Paul Henne and Samual Murray. We want to thank the editors for their extremely helpful comments on an earlier draft of this essay.

In this chapter we explore this tension. We begin in Section 1 by noting that terms like 'consciousness' are ambiguous, and we distinguish between senses that concern notions like being awake, aware, and exercising control and the sense that philosophers have tended to focus on—phenomenal consciousness—suggesting that the former corresponds with the dominant usages outside of academia. We then survey the empirical evidence indicating that non-philosophers do not tend to categorize mental states in the way we would except if they recognize the (supposed) phenomenality of certain mental states. In Section 2, we discuss recent studies of free will attribution that suggest a close connection between these attributions and judgments about consciousness, focusing on work that interprets this in terms of phenomenal consciousness and especially phenomenally conscious emotions. We critically examine this research, arguing that the evidence does not clearly support the conclusion. Finally, in Section 3, we suggest a different picture on which consciousness and free will are treated as more biologically basic features of living animals. We then report the results of a new study building off Ozdemir's (2021) work on the dimensions of mind perception.

1. Concepts of Consciousness

The English word 'consciousness' is used in a number of different ways. As Ned Block (1995, 227) famously notes, "the concept of consciousness is a hybrid, or better, a mongrel concept." Similarly, David Chalmers states that "the term 'consciousness' is ambiguous, referring to a number of different phenomena," including that "sometimes it is used synonymously with 'awakeness'" and "sometimes it is closely tied to our ability to focus attention, or to voluntarily control our behavior" (1996, 6). Each of these senses readily suggest a connection to free will, either as a precondition for it (one must be awake to act freely) or more directly (exhibiting voluntary control of behavior presumably being a hallmark of free action). And each sense is suggested by standard dictionary entries. For instance, the full set of definitions for the adjective 'conscious' given by Oxford's free English dictionary reads as follows:

1. Aware of and responding to one's surroundings; awake.

1.1 Having knowledge of something; aware.

1.2 (conscious of) Painfully aware of; sensitive to.

1.3 Concerned with or worried about a particular matter.

1.4 (of an action or feeling) deliberate and intentional.

1.5 (of the mind or a thought) directly perceptible to and under the control of the person concerned.[2]

Here, definition **1** corresponds with awakeness, while the focus on awareness (as well as in **1.1** and **1.2**) is suggestive of attention, and "deliberate and intentional" in **1.4** indicates voluntary action. Further, similar definitions are found for related terms like 'consciously' (e.g., "in a deliberate and intentional way") and 'consciousness' (e.g, "the state of being awake and aware of one's surroundings").[3] And a preliminary corpus investigation is in line with such definitions.[4] Using the word2vec semantic space for non-academic portions of the Corpus of Contemporary American English built by Sytsma et al. (2019), we find that the nearest neighbor—the term that the model suggests has the closest meaning to 'consciousness' based on the contexts in which they are used—is 'awareness,' while other nearby neighbors include 'wakefulness.'[5]

While a full and careful analysis of the ordinary use of terms like 'consciousness' is beyond the scope of the present chapter, we take the preceding to suggest that the dominant usages relate to being awake, aware, and exercising control. Philosophers of mind, however, have most commonly been concerned with another sense of 'consciousness' that does not clearly coincide with the definitions seen above—what is typically referred to as 'phenomenal consciousness.' This is the sense that both Block and Chalmers are concerned with, with each author raising the comparison in part to note the potential for confusion and to urge that we need to be careful in reasoning about consciousness. And, in fact, recent work on psycholinguistic

[2] Accessed through Lexico.com, 3 June 2021, using the US dictionary. Similar sets of definitions are found for other online dictionaries, as noted in the supplemental materials.

[3] From Lexico.com using the US dictionary, accessed 3 June 2021.

[4] Corpus analysis involves the use of large collections of text for assessing some research question. For instance, the Corpus of Contemporary American English contains over one billion words drawn in a balanced ways from a variety of sources, including spoken English, fiction, popular magazines and newspapers, and academic texts. Terms in the corpus are searchable, including by part of speech and proximity to other terms, and provide the context for utterances of interest. Such a corpus offers a cross-section of "words in the wild," as it were, providing a means of assessing linguistic hypotheses outside of an experimental setting and, hence, free of the potential biases such a setting might introduce. The full text for the corpus is also available and can be used to develop computational models of word meaning by looking at the relative distribution of words across the corpus, such as the distributional model employed here. For a discussion of the value of corpus analysis for philosophy, see Bluhm (2016) and Caton (2020). And, for an extended discussion of corpus methods, see McEnery and Wilson (2001). For an introduction to distributional semantic models see Erk (2012) and Turney and Pantel (2010).

[5] The ten nearest neighbors of "consciousness" in decreasing order of similarity as measured by the cosine of the vectors (in parentheses) are: "awareness" (.69), "unconsciousness" (.65), "wakefulness" (.62), "realization" (.59), "psyche" (.59), "transcendence" (.59), "selfhood" (.59), "subconscious" (.59), "emotion" (.58), and "aliveness" (.58).

biases indicates that people are prone to make inferences licensed only by the dominant or most common sense of a term even when it is used in a subordinate sense (Fischer and Engelhardt 2017, 2020; Fischer and Sytsma 2021; Fischer et al. forthcoming). So, indeed, caution is warranted.

What is the (supposed) phenomenon of phenomenal consciousness? While phenomenal consciousness has been a focus of fierce debate in contemporary philosophy of mind, researchers are not in complete agreement about how it should be understood. Nonetheless, the standard idea is that a range of mental states have something important in common: there is "something it is like" (Nagel 1974) to undergo these states, with this being understood in terms of their having certain introspectively accessible qualities ('phenomenal qualities' or 'qualia' for short), such as the redness we're acquainted with in looking at a ripe tomato. The phrase 'phenomenal consciousness' is, thus, generally taken to characterize a diverse range of mental states as having something important in common—each having phenomenal qualities—and this feature is taken to distinguish these mental states from others that lack such qualities. While there is disagreement about exactly what range of states have phenomenal qualities, they are prototypically taken to include perceptual states and bodily sensations and to exclude at least some intentional states. As Chalmers puts it (2018, 7), "it is widely accepted that seeing a bright red square and feeling pain are phenomenally conscious, and that one's ordinary background beliefs (my belief that Paris is in France, say, when I am not thinking about the matter) are not."

1.1 Attributions of Consciousness

It is often assumed that while phrases like 'phenomenal consciousness' and terms like 'qualia' are clearly technical, the underlying concepts are nonetheless widespread among non-philosophers, and this assumption is made by both skeptics and realists. For instance, Dennett (2005, 27) suggests that belief in qualia is part of our "folk theory of consciousness." On the other end of the spectrum, qualia realists generally take the occurrence of such qualities to be both introspectively obvious and of clear importance. As Chalmers (1995, 207) puts it, qualia are "the most central and manifest aspect of our mental lives." If this is correct, then we would expect the concept of phenomenal consciousness to play a central role in our folk psychological conception of mind. And, in fact, Chalmers (2018, 13) makes this expectation clear, suggesting that "the central intuitions [about phenomenal consciousness] are widely shared well beyond

philosophy" and contending that "it is highly plausible that versions of many of these intuitions can be teased out of ordinary subjects." While Chalmers is focused on what he terms "problem intuitions" here—intuitions reflecting that we take there to be a problem of phenomenal consciousness, including judgments that it is distinctively hard to explain or is non-physical—such intuitions presuppose a conception of phenomenal consciousness that is suitably close to the philosophical concept.[6]

Further, there is some supporting evidence for such claims about the folk theory of consciousness. Knobe and Prinz (2008) present results suggesting that non-philosophers employ the concept of phenomenal consciousness. Most notably, in their second study they gave participants a series of ten statements attributing mental states to a corporation—five that they classified as phenomenal and five that they classified as non-phenomenal—and asked them to rate how weird sounding the statements were. Knobe and Prinz found that on average each of the phenomenal statements was judged to sound less natural than each of the non-phenomenal statements.[7] Another body of evidence, however, suggests that non-philosophers do not tend to employ a concept that is relevantly akin to the philosophical one. Evidence for this was first put forward by Sytsma and Machery (2010), and as detailed below, their contention has since gained further support from a range of follow-up studies, as well as work on beliefs about perception and studies on the dimensions of mind perception.

Sytsma and Machery (2010) argue that if the existence of phenomenal consciousness is pretheoretically obvious and is a central part of our folk psychological conception of mental life, then we would expect the distinction between phenomenal and non-phenomenal states to be clear in how people categorize mental states. To test this hypothesis, in their first study Sytsma and Machery gave philosophers and non-philosophers a description of either a simple, non-humanoid robot or an ordinary human performing one of two simple tasks. The tasks were designed to

[6] We'll focus on the prior question of whether non-philosophers tend to possess such a conception of phenomenal consciousness. If they do not, of course, then it will also follow that they lack further problem intuitions *about* phenomenal consciousness. Nonetheless, assuming for the sake of argument that people do tend to have such a conception, recent empirical work suggests that the majority do not harbor problem intuitions (Peressini 2014, Gottlieb and Lombrozo 2018, Fischer and Sytsma 2021, Ozdemir 2021, Diaz forthcoming).

[7] A number of notable criticisms have been raised against Knobe and Prinz's interpretation of these results, however, including Sytsma and Machery (2009), Arico (2010), and Phelan et al. (2013). Further evidence for non-philosophers potentially having a concept of phenomenal consciousness is provided by Peressini (2014), although he holds that "folk phenomenality is importantly different from the philosophical sense in that it is grounded in the *physical* as opposed to the metaphysical nature of the experiencer" (883). See Sytsma (2014) for a discussion of these works and Sytsma (ms) for an extended critique of Peressini's studies.

elicit the judgment that the human had underwent a prototypical phenomenal state—either seeing red or feeling pain—and the robot was described as behaving in an analogous manner. Participants were then asked about whether the agent had the relevant state.

As expected, philosophers tended to treat the two states similarly, ascribing both to the ordinary human and denying both of the simple robot. In contrast, participants with little to no training in philosophy did not treat them similarly: although they tended to ascribe both states to the ordinary human, they were split with regard to the simple robot, tending to ascribe one prototypical phenomenal state to it (seeing red), while tending to deny the other (feeling pain). These results have subsequently been replicated and expanded upon (Sytsma and Machery 2012, Sytsma 2012, Sytsma and Ozdemir 2019, Cova et al. 2021, Ozdemir 2021). This includes that key results have been found regardless of how reflective the participants are (Sytsma and Machery 2012 responding to Talbot 2012), when using an expanded range of questions to allow participants to distinguish between merely detecting and truly seeing (Sytsma 2013 responding to Fiala et al. 2013), when using phrasings expected by critics to more clearly indicate phenomenal consciousness ("experience red," "experience pain"; Sytsma and Ozdemir 2019), and on larger cross-cultural samples (Sytsma and Ozdemir 2019, Ozdemir 2021).

1.2 Dimensions of Mind Perception

The results reported by Sytsma and Machery (2010), and supported by subsequent follow-up work, suggest that by and large non-philosophers do not tend to treat two prototypical examples of phenomenal states as having something important in common. Given standard justifications for taking the supposed phenomenon of phenomenal consciousness seriously in the first place—that it is central and manifest—finding that non-philosophers do not tend to treat two prototypical examples of supposed phenomenal states similarly in turn suggests that they do not share the philosophical concept of phenomenal consciousness.

This conclusion gains further support from two sources. First, research suggests that non-philosophers tend to hold a naïve view of both colors and pains, treating these as qualities of extra-cranial objects rather than qualities of mental states (Sytsma 2009, 2010, 2012; Reuter 2011; Reuter et al. 2014, 2019; Kim et al. 2016; Sytsma and Reuter 2017; Reuter and Sytsma

2020; Goldberg et al. forthcoming).[8] Since the standard understanding of phenomenal qualities treats them as qualities of mental states, this work suggests that people do not generally think of colors and pains as being phenomenal. But, given that these are two prototypical examples of phenomenal qualities, this raises further doubts about whether they employ such a concept in the first place.

Second, recent work on the dimensions of mind perception suggests that the way non-philosophers tend to attribute various mental states does not respect the philosophical division between phenomenal and non-phenomenal states (Weisman et al. 2017, Malle 2019, Ozdemir 2021). This work seeks to uncover the conceptual structures that frame how people conceive of mental life by looking at patterns of similarity and difference in the capacities they attribute to a range of entities. If the concept of phenomenal consciousness is a central part of folk psychology, then we would expect people to treat clear examples of phenomenal states as tending to hang together. For instance, like the philosophers surveyed by Sytsma and Machery (2010), we would expect them to treat simple robots as being incapable of any phenomenal states, while treating ordinary humans as being capable of the full range of typical phenomenal states. But this is not what we find: across a range of studies, prototypical examples of phenomenal states are found to fall across different dimensions.

While the pioneering work on dimensions of mind perception by Gray et al. (2007) suggests that there are two distinct dimensions, subsequent work often suggests three or more dimensions, although the exact character of these differs somewhat. Weisman et al. (2017) conducted a series of four studies asking people about whether a range of entities—such as a stapler, robot, beetle, adult human—were capable of each of 40 mental capacities. Across these studies, they found a consistent three-dimensional structure to participants' attributions, which they termed *Body*, *Heart*, and *Mind*. Most importantly for present purposes, while the capacities tested included many that philosophers would typically classify as phenomenal, these were spread across the three dimensions. To illustrate, the Body dimension, which Weisman and colleagues characterize in terms of "physiological sensations related to biological needs, as well as the kinds of self-initiated behavior needed to pursue these needs" (11375), included such "phenomenal" capacities as experiencing pain and experiencing pleasure. The Heart dimension

[8] But see Borg et al. (2019), Salomons et al. (forthcoming), and Liu (forthcoming) for responses that suggest a more complicated picture for pains.

was characterized in terms of "basic and social emotions, as well as the kinds of social-cognitive and self-regulatory abilities required of a social partner and moral agent" (11375), and included feeling embarrassed. And the Mind dimension, characterized in terms of "perceptual-cognitive abilities to detect and use information about the environment" (11375), included sensing temperatures and seeing things. Weisman et al.'s goal here was to use patterns of attributions to the different entities "to infer which mental capacities were seen as related and which were considered independent" (11375). Focusing on phenomenal states, their results then suggest that people tend to consider some prototypical examples as being independent of one another.

Malle (2019) also arrived at and replicated a three-dimensional structure across a series of five studies involving attributions of a different set of capacities to a variety of entities. While the dimensions he found are largely consistent with those identified by Weisman et al., they differ in some details and were given different labels—*Affect* (including such prototypically phenomenal capacities as "feeling pain" and "feeling happy"), *Moral & Social Cognition* (including "feel shame or pride" in Study 1), and *Reality Interaction* (including "seeing and hearing the world"). The key point for present purposes, though, is that we once again find that capacities that philosophers would typically classify as phenomenal fell across different dimensions. Insofar as these dimensions capture "people's folk conception of [how] the mind is organized" (2268), such findings again suggest that non-philosophers do not tend to divide up the mental world in a way that follows the philosophical concept of phenomenal consciousness.

Finally, Ozdemir (2021) ran two studies exploring the dimensions of mind perception from an angle more inline with the philosophical discussions and thought experiments. In similar fashion to the studies on free will from Shepherd (2015) and Nahmias et al. (2020) that we will discuss below, participants were given a description of a highly sophisticated robot:

> Imagine that in the future scientists are able to exactly scan a person's brain at the molecular level. Using the information from the brain scan they can create a perfect computer simulation of the working brain. They can then embed that computer in a robot body to create a robot version of the person.
>
> Imagine that scientists scan your brain and use that information to create a robot version of you. On a scale of 0 (Not at all capable) to 6 (Highly capable), how capable would the robot be of…

After reading the description, in Ozdemir's first study participants were asked to rate the same 40 capacities tested by Weisman and colleagues. He found a four-dimensional structure, similar to

those seen previously except for separate dimensions for *Negative States* (including experiencing pain and experiencing fear) and *Positive States* (including experiencing pleasure and feeling love). The other two dimensions were *Perception and Cognition* (including seeing things and sensing temperature) and *Consciousness and Agency* (which included no prototypically phenomenal capacities but did include being conscious). The capacities tested by Weisman and colleagues weren't specifically designed with the philosophical distinction between phenomenal and non-phenomenal states in mind, however. As such, it might be worried that some of the items wouldn't clearly suggest phenomenal states even if participants employ such a notion. Most notably "seeing things" might be interpreted in a purely informational sense along the lines of "detecting objects" rather than a phenomenal sense. Because of this, in his second study, Ozdemir used the same vignette as above, but included two statements involving seeing ("the robot would see colours," "the robot would experience sights and sounds"). Against the worry, and in line with Sytsma and Ozdemir's (2019) extension of Sytsma and Machery's (2010) results, ratings for the 'see' and 'experience' statements were similar (M=6.07, M=5.75), while each was significantly different from ratings for the 'pain' statement (M=3.79); further, while 'see' and 'experience' fell under the same dimension (*Perception and Cognition*), 'pain' was in a different dimension (*Consciousness and Agency*).

To summarize, we saw above that many philosophers hold that the concept of phenomenal consciousness is a central facet of the folk psychological conception of mind, with the existence of phenomenal consciousness being taken to be pretheoretically obvious and its occurrence of central importance to people, both philosophers and non-philosophers alike. If this is the case, then we would expect people to tend to treat at least prototypical examples of phenomenal states similarly, tending to deny them of entities that intuitively lack phenomenal consciousness and to ascribe them to entities that intuitively possess phenomenal consciousness. But this is not what the data suggests. Across a range of studies conducted by different researchers and using different materials and methods, we consistently find that prototypical examples of phenomenal states are treated differently; in fact, they seem to fall across different aspects of the folk conceptualization of mental life. This suggests, minimally, that the concept of phenomenal consciousness does not play a central role in the folk psychological conception of mind. Absent this, however, it is unclear why we should suppose that non-philosophers possess something like the philosophical concept of phenomenal consciousness in the first place. It of

course remains *possible* that they possess such a concept, even if they don't systematically employ it when we would expect them to. Nonetheless, given that non-philosophers by and large don't seem to carve up the mental world in the way we would expect if they were calling on the concept of phenomenal consciousness, we contend that we have defeasible reason for skepticism about such a concept being part of folk psychology.

2. Free Will

While the empirical evidence surveyed in the previous section suggests that the concept of phenomenal consciousness is not part of folk psychology, another prominent body of work in experimental philosophy seems to paint a different picture. A large body of research has sought to understand the ordinary concept of free will, and some of this work suggests that it bears a close relationship to judgments about consciousness, with some interpreting this in terms of phenomenal consciousness (e.g., Shepherd 2012, 2015; Nahmias et al. 2020). In this section we'll survey this research, arguing that while it does indeed indicate a connection between the ordinary concept of free will and *some* notion of consciousness, it is at best unclear that that notion is phenomenal consciousness.

Consider the work of Shepherd (2012). He reports the results of three studies asking participants about a vignette that states that neuroscientists are convinced either that *conscious* or *unconscious* processes direct people's actions, with the exact wording being varied across the studies. This manipulation had a sizable effect on judgments about free will in each study, with ratings being significantly higher in the conscious conditions than the unconscious conditions. Shepherd concludes from this that "consciousness plays a central role in folk conceptions of free will" (926). It is not clear exactly what sense of 'consciousness' he has in mind here, however, and while the discussion highlights awareness it also hints at phenomenal consciousness. For instance, Shepherd states that "the common conception of consciousness involves awareness of the world and ourselves," although he then proceeds to offer illustrations that are typical of discussions of phenomenal consciousness, such as "perceptions of vivid red" and "the feel of anger" (915). The manipulations in the studies are similarly ambiguous, although each emphasizes awareness. Thus, the key notion in the conscious conditions in Studies 1 and 2—"conscious brain events"—is clarified as "brain events within that person's awareness" (917), while Study 3 focuses on "conscious thoughts and events—thoughts and events within an agent's

awareness" (924). Given this focus, Shepherd's studies do not allow us to conclude that judgments about *phenomenal* consciousness, specifically, impacted participants' free will judgments.

Shepherd (2015) extends these results, reporting three further studies involving consciousness manipulations, and again concludes that "many laypeople regard consciousness as important to free will" (944). We contend that as with the previous studies, it is unclear how participants understood the manipulations, a point that Shepherd notes in Section 5. Thus, while he takes his studies to "emphasize elements of *phenomenal* consciousness," he also notes that they "emphasize elements such as deliberation" (942). We're skeptical that this is the case for Shepherd's first two studies, although it is accurate for his third. In the latter case, however, we worry that despite noting elements that philosophers will tend to interpret as emphasizing phenomenal consciousness, non-philosophers will nonetheless focus on deliberation.

The key description in Shepherd's first study is that the agent at issue (Jim) is said to have "no conscious awareness of the right side of his visual field" (932), with the relevant manipulation then involving whether Jim punched someone on his left side (such that "Jim was consciously aware that he was there") or on his right side (such that "Jim had no conscious awareness that he was there"). As in the previous studies, this manipulation had a strong effect on free will judgments, with judgments being significantly higher in the conscious conditions. Shepherd's second study was similar, but now Shepherd describes Jim as having "no conscious control over his left hand" and then punching a person with either his left or his right hand. Despite this change, the manipulation had a comparable effect on free will judgments.

In each of Shepherd's first two studies, while the vignettes *involve* an agent acting on visual information, what they *emphasize* is conscious awareness (Study 1) or conscious control (Study 2). In contrast, Shepherd's third study clearly emphasizes elements that philosophers would tend to interpret in terms of phenomenal consciousness. Participants were given a vignette describing a sophisticated robot that was then said to be either conscious or nonconscious:

> In the future, humans develop the technology to construct humanoid machines. These machines have very sophisticated computers instead of brains, and very intricate movement-generation systems instead of bones, ligaments and muscles. In fact, they are so sophisticated that they look, talk, and act just like humans, and they integrate into human society with no problem at all. The only way to tell if they are a humanoid machine instead of a human being is to look inside of them (by x-ray, for example).

> **Conscious:** These creations are behaviorally just like human beings, and in addition, these creations possess consciousness. They *actually feel* pain, *experience* emotions, *see* colors, and *consciously* deliberate about what to do.
>
> **Nonconscious:** These creations are behaviorally just like human beings. But, these creations do not possess consciousness. They do not *actually feel* pain (even when they say 'Ouch!'), they do not *experience* emotions, they do not *see* colors, and they do not *consciously* deliberate about what to do.

After this, a particular robot was described as performing either a good or a bad action and participants were asked a series of questions, including whether the robot acted of his own free will. Overall, Shepherd found that while participants tended to attribute free will to the conscious robot (M=4.98 on a 1-6 scale), judgments about the nonconscious robot were below midpoint (M=2.97).

As Shepherd suggests, philosophers are likely to read "*actually feel* pain, *experience* emotions, *see* colors" in terms of phenomenal states. This is not the only interpretation available, however. Alternatively, it seems plausible to see this as drawing a distinction between actually undergoing such states and faking it—the robot merely *acting as if* it had those states—which is arguably neutral with regard to phenomenality. Further, as Shepherd notes, the concluding statement that the robot does or does not "*consciously* deliberate about what to do" can be readily understood in terms of another sense of consciousness altogether—one that emphasizes that the action was under the robot's control. Thus, while Shepherd's studies provide further support for the contention that non-philosophers tend to treat consciousness as a requirement for free will, they do not provide strong evidence that they tend to treat *phenomenal* consciousness as a requirement.[9]

Nahmias et al. (2020) have recently expanded on Shepherd's work, reporting the results of two studies involving sophisticated robots. Most importantly, in their second study they gave participants a vignette based on the one used in Shepherd's (2015) third study, but with revised descriptions for the consciousness manipulation that excluded the potentially problematic phrasing noted above ("*consciously* deliberate"). The revised paragraphs read as follows:

> **Conscious:** Furthermore, the robots are able to behave just like human beings, and they also have components that enable conscious experiences. The robots *actually feel* pain,

[9] More recently, Shepherd (2017) finds an effect for consciousness on free will judgments about a person, with the consciousness manipulation emphasizing both control and awareness. And, in a fascinating series of studies, Björnsson and Shepherd (2020) reverse things, showing that whether a "humanoid machine" is described as operating deterministically or indeterministally affects judgements about whether it has a conscious mental life.

> *see* colors, and *experience* emotions. They do not *just appear* to be conscious when they carry out the same behaviors as humans.
>
> **Nonconscious:** Furthermore, the robots are able to behave just like human beings even though they do not have conscious experiences. They have components that process information such that they can carry out all the same behaviors as humans in just the same ways, but when they do so, they *just appear* to feel pain, *just appear* to see colors, and *just appear* to experience emotions.

After reading one of the vignettes, participants were asked a series of questions involving items on free will, basic emotions (e.g., that the robot can feel happiness, anger, etc.), Strawsonian emotions (e.g., that the robot can feel guilt, pride, etc.), and conscious sensations, among others, answering each using a 7-point scale where 1 indicates strong disagreement and 7 indicates strong agreement.

Crucially, Nahmias and colleagues found that the consciousness manipulation had a significant effect on free will judgements. While this is in line with what Shepherd found in his third study, the difference in the mean responses was quite notably smaller than Shepherd reports. While Shepherd found a (very) large effect, reporting a partial eta squared of .34 for the consciousness manipulation, Nahmias et al. found a small effect, with a partial eta squared of just 0.014. Unlike Shepherd, they found that participants did not tend to ascribe free will to the robot in *either* condition, with the means for their free will scale falling below the neutral point in both the conscious condition (M=3.84) and the unconscious condition (M=3.53).[10] Finally, Nahmias and colleagues report the results of a multiple mediation analysis indicating that judgements about Strawsonian and basic emotions fully mediate the effect of the consciousness manipulation on free will judgments, while sensations did not play a significant role.

Nahmias et al. take these results to support their contention that there is a "strong intuition among most people… that the capacity to have conscious experiences is crucial for free will" (61). Their aim, however, is to go beyond this and offer an indication of *why* people treat phenomenal consciousness as being necessary for free will. Nahmias and colleagues suggest that the reason is that people hold that for an agent to have free will things have to *really* matter to the agent, and that this requires phenomenal consciousness. They then take the results of their

[10] Thanks to Eddy Nahmias and Corey Allen for providing these numbers.

mediation analysis to suggest that people take the capacity for phenomenally conscious emotions to be crucial for having things truly matter to an agent.[11]

In our view, the consciousness manipulation used by Nahmias and colleagues is notably clearer than previous attempts, avoiding explicit mention of awareness, deliberation, or control. For a philosophical reader, we have no doubt that the descriptions unambiguously convey that they are talking about phenomenal consciousness. We are less convinced, however, that this will be so clear for non-philosophers.[12] First, as noted in our discussion of the third study in Shepherd (2015), one potential issue is that the descriptions intended to flag phenomenal consciousness might equally be interpreted as drawing a distinction between actually undergoing such episodes and merely faking it. And the added emphasis on whether the robots "*just appear*" to be conscious in Nahmias et al.'s vignettes would seem to further promote such a reading. Second, unlike in Shepherd's study, Nahmias and colleagues report a quite small difference in free will ratings between the two conditions. Because of this, even if their descriptions worked as intended, we do not find that this result offers strong support for the claim that people *generally* treat phenomenal consciousness as being necessary for free will and believe that caution is warranted at this point.

In fact, re-considering Shepherd's results in light of those reported by Nahmias et al., we find that they cast significant doubt on the hypothesis that Shepherd's results reflect judgments about phenomenal consciousness. While there are a number of differences between the two studies, including the exact phrasing of the vignettes and the questions, Nahmias and colleagues based their materials off of Shepherd's and the phrasing is generally similar. The difference that stands out, however, is that Nahmias et al. only emphasized the "phenomenal elements" in the consciousness paragraphs, excluding the description of whether the robot could or could not consciously deliberate about what to do. Thus, one plausible hypothesis is that the much larger effect found by Shepherd primarily reflects judgments about conscious deliberation, *not* phenomenal consciousness.

---

[11] Note that the proposed conceptual connection between consciousness and emotions gains some support from the nearest neighbors analysis in Section 1, where "emotion" was one of the ten closest terms to "consciousness" in the semantic space.

[12] Arguably in line with this, although it was stipulated that the robot has conscious experiences in the conscious condition, the mean responses for the sensation items—e.g., "These robots *experience*, more than just *process*, the sounds in music, the images in art, the smells of food, and the softness of a blanket." (70)—was below the neutral point. Similarly, despite having specified that the robots *actually experience* emotions, mean ratings for both basic emotions and Stawsonian emotions were below the midpoint.

3. Further Explorations

While the studies discussed above provide considerable evidence that non-philosophers take consciousness *in some sense* to be related to free will, we find it to be far less clear that these findings specifically reflect judgments about *phenomenal* consciousness. Instead, we suspect that participants have tended to interpret the consciousness talk in these studies in ways that concern being awake, aware, and exercising control. We believe that this alternative explanation is quite plausible. First, as noted in Section 1, this cluster of concepts bear a clear relation to free action, such that the results are readily explicable on this interpretation. Second, we saw preliminary evidence that this cluster corresponds with the dominant ordinary usages of terms like 'consciousness' outside of academia. Third, we saw reason to doubt that non-philosophers tend to employ the philosophical concept of phenomenal consciousness. Thus, absent strong evidence to the contrary, we believe the *default assumption* should be that people will interpret consciousness talk in terms of being awake, aware, and exercising control rather than phenomenal consciousness. And, fourth, looking across the studies from Shepherd (2015) and Nahmias et al. (2020) discussed above, the effect of the consciousness manipulation on free will judgments was greatly diminished when reference to 'conscious deliberation' was excluded.

In the remainder of this chapter, we'll further explore the relationship between consciousness and free will. We begin by returning to the work on the dimensions of mind perception discussed above in light of Nahmias et al.'s research, now with a focus on judgments about free will and emotions. We then add to this body of work.

3.1 Revisiting the Dimension Studies

In Section 1, we discussed recent work on the dimensions of mind perception with a focus on prototypical phenomenal states such as feeling pain and seeing red. Several of these studies also included a question about free will, however, potentially offering insight into how this fits into the folk conception of how the mind is organized.

Perhaps most importantly, Weisman et al. (2017) found that having free will fell within the Body dimension, not the Heart dimension as Nahmias et al. might plausibly predict. As Weisman et al. put it, "having free will tracked the physiological phenomena of the body more closely than the social-emotional capacities of the heart" (11376). We saw above that Nahmias et al.'s basic picture is that one crucial aspect of thinking about an agent having free will is that

things can truly matter to the agent, and they suggest that this in turn requires phenomenal consciousness—that the agent can experience the consequences of her actions—focusing especially on the ability to feel emotions. But Weisman et al.'s work doesn't fit cleanly with this picture. As noted above, the Body dimension is characterized in terms of bodily sensations and capacities related to meeting biological needs. And while some basic emotions fall under this dimension, they relate especially to biological safety (e.g., experiencing fear, feeling safe). In contrast, most emotions fell under Weisman et al.'s Heart dimension. Recalling that these dimensions reflect the items that tended to be treated as going together across a range of entities, what we take this to suggest is that people tend to treat free will as being a more biologically basic feature of living animals than social emotions.[13]

Interestingly, Weisman et al. also found that being conscious and being self-aware fell under the Body dimensions. This is plausibly in keeping with the analysis of 'consciousness' in Section 1 and the alternative explanation we offered of the effect of consciousness manipulations on free will judgments. It suggests a basic biological conception of consciousness that is congruent with a focus on being awake, aware, and exercising control. We might think of these as "abilities related to the physical, biological body" of living animals, as Weisman et al. (11375) describe the Body dimension.[14]

While the results for Ozdemir's (2021) dimension studies are largely congruent with Weisman et al.'s findings, they paint a less clear picture with regard to having free will. In his first study, this item did not load especially strongly on any of the dimensions he identified, while the various emotions tested grouped together under either the Negative State or Positive State dimensions. In Ozdemir's second study, however, both having free will and having emotions fell within his broad Consciousness and Agency dimension, as did items like being alive and being aware of oneself. While this is congruent with the picture drawn from Weisman et al., it is also compatible with Nahmias et al.'s contention that free will is more closely tied to feeling emotions. As such, we conclude by expanding upon Ozdemir's second study, reporting

[13] Compare, for example, the neutral to positive mean ratings for having free will for a beetle, goat, and elephant to the low mean ratings for feeling embarrassed and experiencing guilt for these animals in Weisman et al.'s Figure 1.
[14] Recall that "aliveness" was one of the nearest neighbors for "consciousness" in the analysis above. See Arico et al. (2011) for experimental evidence that judgments that something is alive might be a basic cue for a range of mental state attributions.

the results of a new study based on his materials and employing an analysis aimed at revealing the relations between the items comprising his broad Consciousness and Agency dimension.

3.2 Clustering Attributions

To further explore the relationship between the items in Ozdemir's (2021) second study, we tested two variations on his vignette. We then employed bottom-up cluster analyses to help assess the patterns of judgments across the 25 capacities he tested.

As detailed above, the vignette for Ozdemir's study involved a person's brain being scanned to create a computer simulation that was then embedded in a robot body. We updated this to further emphasize that the resulting being was a *physical* duplicate of the original person, doing so in two different ways. The first variation stuck closely to Ozdemir's original scenario, but specified that both the person's brain *and* body were scanned to build an android duplicate. The second variation went further, replacing the android with a "molecule by molecule" biological duplicate reminiscent of philosophical zombies (Chalmers 1996), and specifying that it was only the physical make-up of the person that was duplicated, such that "if there was any non-physical aspect to the person—such as a non-physical soul or mind—the scientists would not be able to duplicate that aspect of the person." Participants were given one of the two vignettes, then asked whether they agreed or disagreed with each of the 25 statements used by Ozdemir (with "duplicate" replacing "robot"), in random order, answering on a 1–7 scale where 1 indicates "Disagree Strongly" and 7 indicates "Agree Strongly." In addition, they were asked a range of demographic and check questions. Participants were recruited globally through advertising for a free personality test on Google Ads, with the ad and questionnaire presented in English. Responses were collected from 886 participants who were at least 16 years of age and completed the philosophical questions. The full text for both vignettes, as well as extended details about the questions, sample, and analysis are provided in the supplemental materials.[15] For space, we'll focus on giving just a high-level overview of the findings here.[16] Results are shown in Figure 1.

[15] http://philsci-archive.pitt.edu/19556

[16] Given the diversity of the sample exploratory tests were run to check for demographic differences. As few notable differences were found, we'll focus on the full sample here.

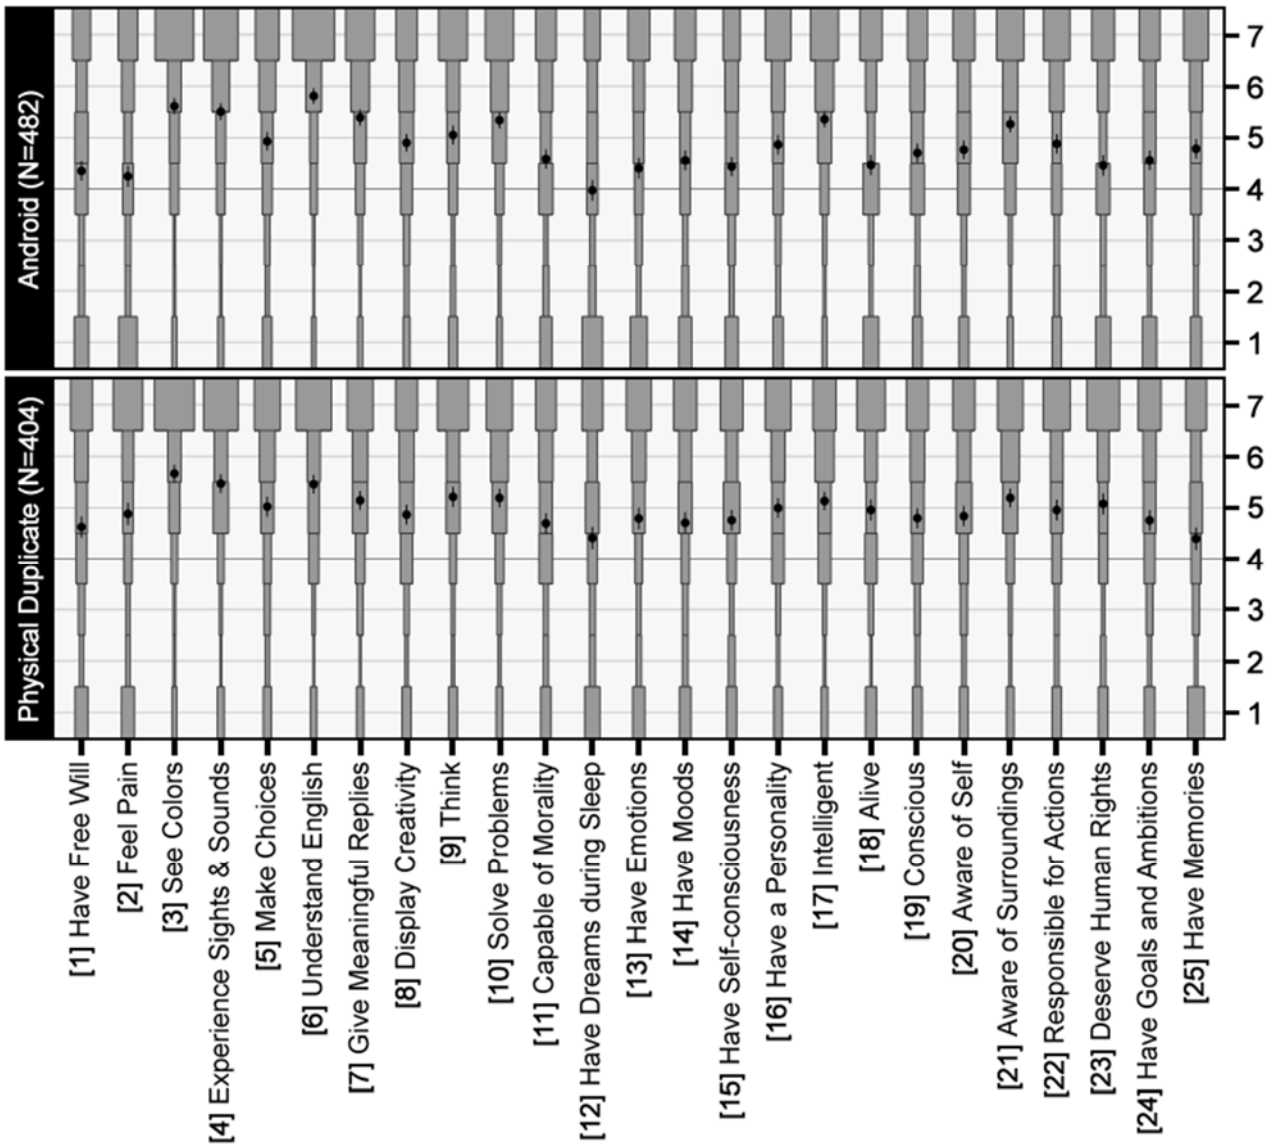

Figure 1: Study results broken down by condition, with the relative percentage of participants selecting each response option shown in gray and the mean and 95% confidence interval overlaid in black.

To assess differences between the conditions, Welch's t-tests were run for each of the 25 items, applying the Holm–Bonferroni correction. Only three showed a significant difference, and even here the effect sizes were small: ratings for **[2]** *Feel Pain*, **[18]** *Alive*, and **[23]** *Deserve Human Rights* were significantly higher for the Physical Duplicate (M=4.88, 4.96, 5.07) than the

Android (*M*=4.25, 4.47, 4.46).[17] To assess general agreement or disagreement with the items, we ran one-sample Student's t-tests with the Holm–Bonferroni correction for each condition. Each item was above and significantly different from the midpoint, with the exception of **[12]** *Have Dreams during Sleep* for the Android. Thus, interestingly, participants tended to treat the android as being alive, although there was disagreement on this score. And the general agreement found for **[19]** *Conscious* casts further doubt on the claim that "problem intuitions" are widespread (Chalmers 2018), and specifically intuitions that consciousness is non-physical (see Footnote 7).

For present purposes, though, our primary concern is with the relation between different judgments about the duplicates. To assess this, we began by calculating distance matrixes for the items using two distance measures (Euclidean, Spearman). These provide measures of how similar participants' responses were between each pair of items. Euclidean, or straight-line distance, is perhaps the most common in hierarchical clustering and is often treated as the default. This method treats the ratings as having interval scale, however, and while this is a common assumption for Likert-ratings, it might not hold. The Spearman distance measure, by contrast, is a rank-based correlation metric. As detailed in the supplemental materials, the results were similar for each measure. This includes that pairs of items that we expected to be treated similarly were in fact close together on both measures, offering some indication that participants understood the task: **[3]** *See Colors* was the closest item to **[4]** *Experience Sights & Sounds* for each, and vice versa, with this being the closest pair overall by Euclidean distance; and likewise for **[13]** *Have Emotions* and **[14]** *Have Moods*, which were the closest pair overall by Spearman distance. By contrast, **[1]** *Have Free Will* was not so notably close to **[13]** *Have Emotions* on either measure, with **[14]** *Have Moods*, **[15]** *Have Self-consciousness*, **[16]** *Have a Personality*, and **[12]** *Have Dreams during Sleep* being closer on each (and **[23]** *Deserve Human Rights*, **[9]** *Think*, **[2]** *Feel Pain*, and **[11]** *Capable of Morality* also being closer using Spearman distance).

[17] The result for **[2]** *Feel Pain* is in keeping with previous work indicating that what an agent's body is made of impacts people's pain judgments (Sytsma 2012). Similarly, the results for the other two statements seem readily explicable in terms of difference between a robotic and an organic system.

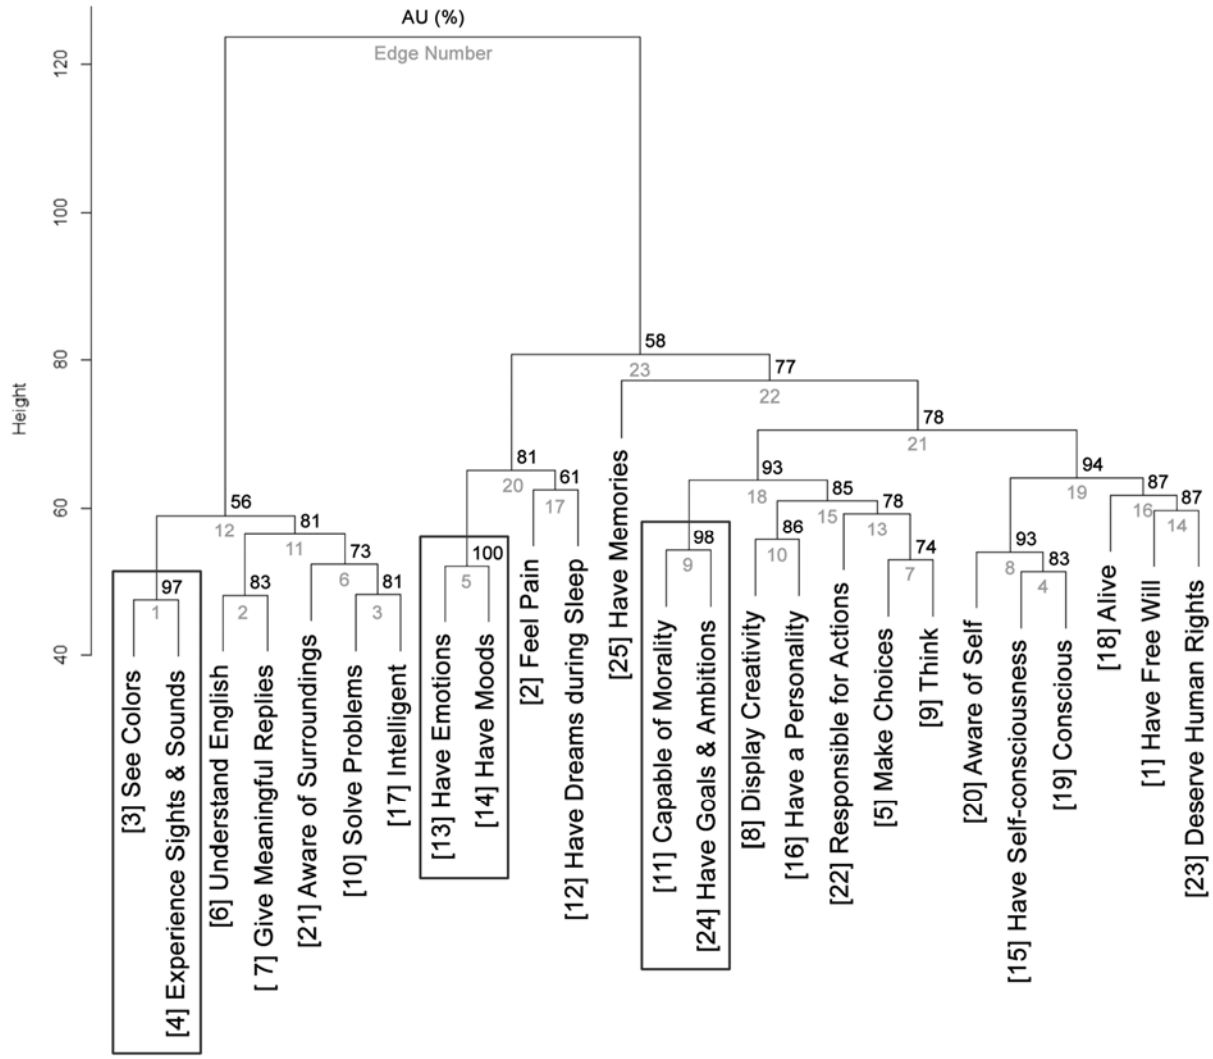

Figure 2: Dendrogram for the hierarchical cluster analysis using Ward's method with the Euclidian distance measure. Edge numbers are displayed below (grey) and AU values are displayed above (black). Highest-level clusters occurring in 95% or more of the resamples are highlighted. Item numbers correspond with those given in Figure 1.

To better assess the patterns of similarity between item ratings, we performed a series of agglomerative hierarchical cluster analyses. This works in a bottom-up manner based on the distance matrix, starting with each item being considered as a single-element cluster, then iteratively combining the most 'similar' clusters until all items are grouped together. By working in a bottom-up manner, agglomerative clustering tends to do a better job of identifying smaller clusters, making it a good choice for present purposes. The specific clustering produced depends not only on the distance measure, but on the clustering method employed. We tested five

clustering methods for each of the two distance measures. These were then assessed with two goals in mind—faithfulness to the distance matrix and the strength of clustering produced. Based on this, we identified the two best methods for each measure. This included what is perhaps the most common combination—Ward's method with Euclidean distance. We'll focus on this clustering here, but the other three analyses paint a similar picture, as detailed in the supplemental materials. To assess the uncertainty of the analyses we used a resampling method with 10k iterations to produce approximately unbiased (AU) p-values. This gives the percentage of iterations where the same cluster was found, and we highlight clusters with values of 95 or greater in the dendrograms. The dendrogram for Ward's method with Euclidean distance is shown in Figure 2.

The first thing to note is that the two high-level clusters in the dendrogram largely correspond with the dimensions found by Ozdemir (with the exception of **[5]** *Make Choices* and **[8]** *Display Creativity*), while the lowest-level clusters include the pairs of items discussed above that were expected to be similar (**[3]** *See Colors* and **[4]** *Experience Sights & Sounds*, **[13]** *Have Emotions* and **[14]** *Have Moods*). This offers some validation for the clustering. Further, in line with Ozdemir's findings, we again find that the prototypical examples of phenomenal states—**[2]** *Feel Pain* and **[3]** *See Colors* / **[4]** *Experience Sights & Sounds*—fell across different top-level clusters. This adds to the findings discussed in Section 1, lending further support to Sytsma and Machery's (2010) conclusion.

Most importantly for present purposes, looking more closely at the structure produced, we find that **[1]** *Have Free Will* falls in a mid-level cluster with **[18]** *Alive*, **[23]** *Deserve Human Rights*, **[19]** *Conscious*, **[15]** *Have Self-consciousness*, and **[20]** *Aware of Self*. This is in line with the expected connection between free will and consciousness, further suggesting that people treat these as going together. In contrast, while **[1]** *Have Free Will* falls within the same top-level cluster as **[13]** *Have Emotions*, they weren't otherwise especially close, with **[13]** *Have Emotions* grouping together with the bodily state of **[2]** *Feel Pain* as well as **[12]** *Have Dreams during Sleep*. This is largely congruent with the picture drawn from Weisman et al.'s work. The structure again suggests that people treat free will as being closely connected to consciousness and self-awareness, with these being tied to being alive and moral patiency. While we consider these results to best be thought of as exploratory, we also find them to be congruent with the picture we've suggested: Even focusing on sci-fi examples reminiscent of philosophical thought

experiments, we find that judgments about free will, consciousness, and self-awareness are closely connected to judgments about being alive, suggesting that these form a basic cluster in people's thinking that is congruent with thinking about consciousness in terms of being awake, aware, and exercising control.

4. Conclusion

A large body of empirical evidence indicates that the judgments of non-philosophers about consciousness and free will are related. And some have interpreted the former in terms of the philosophical concept of phenomenal consciousness, with Nahmias et al. (2020) suggesting that people treat the capacity for phenomenally conscious emotions as critical for having free will. Several other bodies of research, however, at the very least indicate that caution is warranted with regard to interpreting non-philosophers' understanding of consciousness in this way. Further, standard dictionary definitions and a preliminary corpus analysis support this conclusion. A critical look at several key empirical studies casts doubt on their providing evidence for there being a close relationship between non-philosophers' judgments about free will and *phenomenal* consciousness. This skeptical conclusion is further supported by the results a new exploratory study, which suggests that attributions of free will are not so closely tied to attributions of emotions. Rather we suggest that both free will and consciousness are treated as being more biologically basic than this.

References

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(3): 327–352.

Björnsson, Gunnar and Joshua Shepherd (2020). "Determinism and attributions of consciousness." *Philosophical Psychology*, 33(4): 549–568.

Block, Ned (1995). "On a confusion about a function of consciousness." *Behavioral and Brain Sciences*, 18: 227–287.

Bluhm, Roland (2016). “Corpus analysis in philosophy.” In M. Hinton (ed.), *Evidence, Experiment and Argument in Linguistics and the Philosophy of Language*. Frankfurt am Main: Peter Lang, pp. 91–109.

Borg, Emma, Richard Harrison, James Stazicker, and Tim Salomons (2019). “Is the folk concept of pain polyeidic?” *Mind & Language*, 35: 29–47.

Caton, Jacob (2020). “Using Linguistic Corpora as a Philosophical Tool.” *Metaphilosophy*, 51(1): 51–70.

Chalmers, David (1995). “Facing up to the problem of consciousness.” *Journal of Consciousness Studies*, 2(3): 200–219.

Chalmers, David (1996). *The Conscious Mind*. New York: Oxford University Press.

Chalmers, David (2018). “The meta-problem of consciousness.” *Journal of Consciousness Studies*, 25 (9-10): 6–61.

Cova, Florian et al. (2019). “Estimating the Reproducibility of Experimental Philosophy.” *Review of Philosophy and Psychology*, 12: 9–44.

Dennett, Daniel (2005). *Sweet Dreams*. Cambridge: MIT Press.

Diaz, Rodrigo (forthcoming). “Do people think consciousness poses a hard problem? Empirical evidence on the meta-problem of consciousness.” *Journal of Consciousness Studies*.

Erk, Katrin (2012). “Vector space models of word meaning and phrase meaning: A survey.” *Language and Linguistics Compass*, 6(10): 635–653.

Fiala, Brian, Adam Arico, and Shaun Nichols (2013). “You, robot.” In E. Machery and E. O’Neill (eds.), *Current Controversies in Experimental Philosophy*, New York: Routledge.

Fischer, Eugen and Paul Engelhardt (2017). “Diagnostic Experimental Philosophy.” *Teorema: International Journal of Philosophy*, 36(3): 117–137.

Fischer, Eugen and Paul Engelhardt (2020). “Lingering stereotypes: Salience bias in philosophical argument.” *Mind and Language*, 35: 415–439.

Fischer, Eugen, Paul Engelhardt, and Justin Sytsma (forthcoming). “Inappropriate stereotypical inferences? An adversarial collaboration in experimental ordinary language philosophy.” *Synthese*.

Fischer, Eugen and Justin Sytsma (2021). “Zombie Intuitions.” *Cognition*, 215: 104807.

Goldberg, Benjamin, Kevin Reuter, and Justin Sytsma (forthcoming). "A brief history of pain concepts." In K. Hens and A. De Block (eds.), *Advances in Experimental Philosophy of Medicine*, Bloomsbury.

Gottlieb, Sara and Tania Lombrozo (2018)." Can science explain the human mind? Intuitive judgments about the limits of science." *Psychological Science*, 29: 121–130.

Gray, Heather, Kurt Gray, and Daniel Wegner (2007). "Dimensions of mind perception." *Science*, 315(5812): 619.

Kim, Hyo-eun, Nina Poth, Kevin Reuter, and Justin Sytsma (2016). "Where is your pain? A Cross-cultural Comparison of the Concept of Pain in Americans and South Koreans." *Studia Philosophica Estonica*, 9(1): 136–169.

Liu, Michelle (forthcoming). "The Polysemy View of Pain." *Mind & Language*.

Malle, Bertram (2019). "How many dimensions of mind perception really are there?" In A. Goel, C. Seifert, and C. Freska (eds.), *Proceedings of the 41st Annual Meeting of the Cognitive Science Society*, 2268-2274, Cognitive Science Society.

McEnery, Tony and Andrew Wilson (2001). *Corpus Linguistics, Second Edition*. Edinburgh: Edinburgh University Press.

Murray, Samuel, Elise Murray, and Thomas Nadelhoffer (2021). "Piercing the smoke screen: Dualism, Free Will, and Christianity." *Journal of Cognition and Culture*, 21: 94–111.

Nahmias, Eddy, Corey Allen, and Bradley Loveall (2020). "When do robots have free will? Exploring the relationships between (attributions of) consciousness and free will." In B. Feltz, M. Missal, and A. Sims (eds.), *Free Will, Causality, and Neuroscience*, 57–80, Brill.

Oxford University Press (2021). Definition of conscious [online]. Available at: https://www.lexico.com/en/definition/wake (Accessed: 3 June 2021).

Ozdemir, Eyuphan (2021). *Empirical Evidence Against Phenomenal Theses*. PhD Dissertation, Victoria University of Wellington.

Peressini, Anthony (2014). "Blurring two conceptions of subjective experience: Folk versus philosophical phenomenality." *Philosophical Psychology*, 27(6): 862–889.

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.

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*, Bloomsbury.

Reuter, Kevin, Michael Sienhold, and Justin Sytsma (2019). “Putting Pain in its Proper Place.” *Analysis*, 79(1): 72–82.

Reuter, Kevin and Justin Sytsma (2020). “Unfelt Pain.” *Synthese*, 197: 1777–1801.

Salomons, Tim, Richard Harrison, Nat Hansen, James Stazicker, Astrid Grith Sorensen, Paula Thomas, and Emma Borg (forthcoming). “Is pain ‘all in your mind’? Examining the general public’s views of pain.” *Review of Philosophy and Psychology*.

Shepherd, Joshua (2012). “Free will and consciousness: Experimental studies.” *Consciousness and Cognition*, 21: 915–927.

Shepherd, Joshua (2015). “Consciousness, free will, and moral responsibility: Taking the folk seriously.” *Philosophical Psychology*, 28(7): 929–946.

Shepherd, Joshua (2017). "The Folk Psychological Roots of Free Will." In D. Rose (ed.), *Experimental Metaphysics*, 95–115, Bloomsbury.

Sytsma, Justin (2009). “Phenomenological Obviousness and the New Science of Consciousness.” *Philosophy of Science*, 76(5): 958–969.

Sytsma, Justin (2010). “Dennett’s Theory of the Folk Theory of Consciousness.” *Journal of Consciousness Studies*, 17(3-4): 107–130.

Sytsma, Justin (2012). “Revisiting the valence account.” *Philosophical Topics*, 40(2): 197–198.

Sytsma, Justin (2013). “The Robots of the Dawn of Experimental Philosophy of Mind.” In E. Machery and E. O’Neill (eds.), *Current Controversies in Experimental Philosophy*, Routledge.

Sytsma, Justin (2014). “Attributions of Consciousness.” *WIREs Cognitive Science*, 5: 635–648.

Sytsma, Justin (ms). “Experiencers and the Ambiguity Objection.”

Sytsma, Justin, Roland Bluhm, Pascale Willemsen, and Kevin Reuter (2019). “Causal Attributions and Corpus Analysis.” In E. Fischer and M. Curtis (eds.), *Methodological Advances in Experimental Philosophy*, 209–238, Bloomsbury Press.

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 (2012). “On the Relevance of Folk Intuitions: A Reply to Talbot.” *Consciousness and Cognition*, 21: 654–660.

Sytsma, Justin and Eyuphan Ozdemir (2019). “No Problem: Evidence that the Concept of Phenomenal Consciousness is not Widespread.” *Journal of Consciousness Studies*, 26(9-10): 241–256.

Sytsma, Justin and Kevin Reuter (2017). “Experimental Philosophy of Pain.” *Journal of Indian Council of Philosophical Research*, 34(3): 611–628.

Talbot, Brian (2012). “The irrelevance of folk intuitions to the ‘hard problem’ of consciousness.” *Consciousness and Cognition*, 21(2): 644–650.

Turney, Peter and Patrick Pantel (2010). “From frequency to meaning: Vector space models of semantics.” *Journal of Artificial Intelligence Research*, 37: 141–188.

Weisman, Kara, Carol Dweck, and Ellen Markman (2017). “Rethinking people's conceptions of mental life.” *PNAS*, 114(43): 11374–11379.